



Administering Avaya Identity Engines Ignition Server

Release 9.1
NN47280-600
Issue 05.01
March 2015

© 2015 Avaya Inc.

All Rights Reserved.

Notice

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

"Documentation" means information published by Avaya in varying mediums which may include product information, operating instructions and performance specifications that Avaya may generally make available to users of its products and Hosted Services. Documentation does not include marketing materials. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of documentation unless such modifications, additions, or deletions were performed by Avaya. End User agrees 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 websites referenced within this site or documentation 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 Avaya hardware and software. 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 website: <http://support.avaya.com> or such successor site as designated by Avaya. Please note that if You acquired the product(s) from an authorized Avaya Channel Partner outside of the United States and Canada, the warranty is provided to You by said Avaya Channel Partner and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, [HTTP://SUPPORT.AVAYA.COM/LICENSEINFO](http://support.avaya.com/licenseinfo) OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AVAYA CHANNEL PARTNER. 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 CHANNEL PARTNER; 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").

Avaya grants You a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The

applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to You. "Software" means computer programs in object code, provided by Avaya or an Avaya Channel Partner, whether as stand-alone products, pre-installed on hardware products, and any upgrades, updates, patches, bug fixes, or modified versions thereto. "Designated Processor" means a single stand-alone computing device. "Server" means a Designated Processor that hosts a software application to be accessed by multiple users. "Instance" means a single copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine ("VM") or similar deployment.

License types

Designated System(s) License (DS). End User may install and use each copy or an Instance of the Software only on a number of Designated Processors up to the number indicated in the order. Avaya may require the Designated Processor(s) to be identified in the order by type, serial number, feature key, Instance, location or other specific designation, or to be provided by End User to Avaya through electronic means established by Avaya specifically for this purpose.

Heritage Nortel Software

"Heritage Nortel Software" means the software that was acquired by Avaya as part of its purchase of the Nortel Enterprise Solutions Business in December 2009. The Heritage Nortel Software currently available for license from Avaya is the software contained within the list of Heritage Nortel Products located at <http://support.avaya.com/LicenseInfo> under the link "Heritage Nortel Products" or such successor site as designated by Avaya. For Heritage Nortel Software, Avaya grants You a license to use Heritage Nortel Software provided hereunder solely to the extent of the authorized activation or authorized usage level, solely for the purpose specified in the Documentation, and solely as embedded in, for execution on, or for communication with Avaya equipment. Charges for Heritage Nortel Software may be based on extent of activation or use authorized as specified in an order or invoice.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, Hosted Service, or hardware provided by Avaya. All content on this site, the documentation, Hosted Service, and the product 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 unless expressly authorized by Avaya. 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.

Virtualization

Each product has its own ordering code and license types. Note that each Instance of a product must be separately licensed and ordered. For example, if the end user customer or Avaya Channel Partner would like to install two Instances of the same type of products, then two products of that type must be ordered.

Third Party Components

"Third Party Components" mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements ("Third Party Components"), which contain terms regarding the rights to use certain portions of the Software ("Third Party Terms"). As required, information regarding distributed Linux OS source code (for those products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the products, Documentation or on Avaya's website at: <http://support.avaya.com/Copyright> or such successor site as designated by Avaya. You agree to the Third Party Terms for any such Third Party Components.

Preventing Toll Fraud

"Toll Fraud" is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

If You suspect that You are being victimized by Toll Fraud and You need technical assistance or support, call Technical Service Center Toll Fraud Intervention Hotline at +1-800-643-2353 for the United States and Canada. For additional support telephone numbers, see the Avaya Support website: <http://support.avaya.com> or such successor site as designated by Avaya. Suspected security vulnerabilities with Avaya products should be reported to Avaya by sending mail to: securityalerts@avaya.com.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support website: <http://support.avaya.com>, or such successor site as designated by Avaya.

Contact Avaya Support

See the Avaya Support website: <http://support.avaya.com> for product or Hosted Service notices and articles, or to report a problem with your Avaya product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: <http://support.avaya.com> (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the Documentation, Hosted Service(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, Hosted Service(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. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Contents

Chapter 1: Introduction	17
Purpose.....	17
Related resources.....	17
Documentation.....	17
Training.....	18
Viewing Avaya Mentor videos.....	18
Subscribing to e-notifications.....	18
Searching a documentation collection.....	20
Support.....	21
Chapter 2: New in this release	22
Features.....	22
Mobile Device Management support.....	22
Avaya Identity Engines Ignition Access Portal enhancements.....	22
Extended High Availability.....	23
Guest Manager enhancements.....	23
Advanced RADIUS proxy.....	23
Chapter 3: Introduction to Avaya Identity Engines Ignition Server	25
What is Avaya Identity Engines Ignition Server?.....	25
Key characteristics of Ignition Server.....	26
The Ignition Server approach.....	27
Wide set of criteria for policy decisions.....	27
Consolidation: Efficiency and clear lines of responsibility.....	27
Compliance automation.....	28
Ignition Server feature overview.....	28
Policy and Directory integration features.....	28
Authentication, Authorization, and Accounting Features.....	29
Platform and OS services.....	30
Administration, Control, and Configuration features.....	30
Security features.....	30
Chapter 4: Ignition Dashboard	32
Administration roles.....	32
Launching Ignition Dashboard.....	33
Dashboard best practices and design usage guidelines.....	34
Initial default display.....	35
Connection indicator.....	35
Error indicator.....	35
Dashboard layout.....	36
Multiple logins alert.....	36
Managing multiple Ignition Server sites.....	37

Setting up a Site Group.....	38
Idle time-out.....	40
Reactivating a session.....	40
Starting a new session.....	41
Quitting from a timed-out session.....	41
Session time-out.....	42
Root certificate.....	42
Administration functionality.....	42
Navigating the Dashboard Hosts.....	43
Navigating the Admin Access Policies.....	44
Navigating the Admin Roles.....	46
Navigating the Admin logs.....	49
Configuring administration preferences.....	51
Configuring the idle time-out for Dashboard.....	51
Setting viewing preferences for the Monitor view.....	52
Refreshing the Ignition Dashboard view.....	53
Exiting Ignition Dashboard.....	53
Checking the Dashboard software version.....	53
Chapter 5: Sites, nodes, and settings.....	54
Introduction to Dashboard's configuration view.....	54
Managing a site.....	55
Site actions.....	55
Renaming an Ignition Server site.....	56
Changing the System Administrator login name.....	56
Configuring the System Administrator password.....	57
Managing Ignition Server services.....	57
Configuring Ignition Server's RADIUS service.....	57
Editing RADIUS communication settings.....	58
Configuring Ignition Server's SOAP service.....	59
Resetting the SOAP password.....	61
Managing a node.....	62
Actions menu for a node.....	62
Rebooting a node.....	62
Powering down a standalone node.....	63
Reinitializing Ignition Server from Dashboard.....	63
Viewing logs for a node.....	64
Renaming a node.....	64
Status tab.....	65
Obtaining the Ignition Server Serial Number.....	65
System tab.....	67
Viewing Ignition Server's DNS settings.....	67
Editing Ignition Server's DNS settings.....	68
Setting the Network Routing configuration.....	68

Adding a route to Ignition Server's Routing Table.....	69
Editing an existing route	69
Deleting an existing Route.....	70
SNMP settings.....	70
SSH settings.....	70
SMTP settings.....	71
Configuring the Ignition Server's network ports	71
Port configuration settings.....	71
Configuring the Admin port	72
Configuring a service port	73
Enabling the service port	73
Configuring the HA port	74
Enabling the HA port	74
Managing Ignition Server licenses	75
Checking Ignition Server licences.....	75
Seat limit enforcements	76
Obtaining PLDS licenses.....	77
Obtaining KRS licenses.....	77
Installing an Ignition Server license	78
Replacing an Ignition Server license.....	79
Making a backup copy of your Ignition licenses	79
Transferring a License to a different Ignition server	80
Troubleshoot tab.....	80
Running a ping test.....	80
Running a packet capture	81
Chapter 6: Managing certificates.....	82
Required types of Certificates	82
Sample certificates.....	82
Format of certificate files	83
Certificates tab	83
Getting a new certificate.....	84
Create the certificate request	84
Import the certificate	85
Assign the certificate for use in Ignition Server.....	86
Admin certificate	86
Installing Dashboard's copy of the Admin Certificate	87
Replacing the Admin certificate.....	88
Assigning the SOAP service certificate	90
Assigning protocol credential certificates.....	91
Installing protocol root certificates	92
Adding a certificate revocation list URL	92
Adding a Certificate revocation List URL to Ignition Server	92
Viewing the certificate revocation list URLs	93

CLI Command to toggle CRL check level.....	95
Viewing a certificate	95
Deleting a certificate or certificate request	95
Viewing an existing certificate request	96
Chapter 7: Authenticators	97
Introduction to Authenticators	97
Matching an incoming request to an authenticator record	97
Authenticator hierarchy and containers	98
How you build your hierarchy.....	99
Using the hierarchy in your policies.....	100
Default container	100
Window layout.....	100
Creating the authenticator hierarchy.....	101
Adding a container to an authenticator hierarchy	101
Renaming a container in an authenticator hierarchy.....	102
Deleting a container from an authenticator hierarchy.....	102
Creating an authenticator	104
A note on Authenticator Vendor and Device Template.....	104
Importing authenticators.....	107
Exporting authenticators.....	108
Registering Access Portal with the Ignition Server	108
Placing an authenticator in the authenticator hierarchy.....	110
Finding an authenticator.....	111
Pinging an authenticator	112
RADIUS global authenticator	112
Creating the Global Authenticator.....	113
RADIUS sub-authenticators.....	114
Viewing and editing sub-authenticators	114
Creating a sub-authenticator	115
Sorting sub-authenticators.....	116
Processing authenticator requests.....	117
How Ignition server processes RADIUS requests from authenticator bundles	117
Changing the configuration of an authenticator	117
Removing an authenticator from its place in the hierarchy	118
Renaming authenticators.....	118
Deleting authenticators.....	119
Chapter 8: Using Ignition Access Portal as a Policy Enforcement Point	120
Registering Access Portal with the Ignition Server	121
Configuring a guest access policy.....	123
Chapter 9: Internal users, groups, and devices	127
Ignition Server Internal Data Store	127
Internal users.....	127
Internal Users Panel.....	128

Filtering the internal users list.....	129
Viewing an Internal User	129
Creating an Internal User.....	129
Importing user records.....	132
Exporting user records.....	133
Internal devices	133
Finding an internal device	134
MAC Addresses are stored only in the internal store.....	134
Filtering the device list	134
Creating a device record.....	135
Assigning a device to a user or group.....	137
Editing a device record.....	138
Deleting a device record	138
Importing device records.....	138
Exporting device records	139
MAC address wildcarding	140
Adding an internal device from the Monitoring Access Logs.....	141
Internal groups	142
Internal groups panel	142
Working with the Internal Groups Hierarchy.....	143
About the default user group.....	144
Adding a new internal group	144
Moving an internal group in the hierarchy.....	144
Renaming an internal group	145
Changing a group’s type designation.....	145
Deleting an Internal Group.....	145
Working with internal group details.....	146
Chapter 10: Directory Services.....	148
Quickstart: Directory Services in Dashboard	148
Introduction to Directory Services	148
Directory Services	149
Supported Directory Servers.....	149
Internal Data Store as a Directory Service.....	150
Working With Directory Services.....	150
Connecting to active directory	151
AD connection settings	151
Preparing to connect to Active Directory.....	153
Creating an Active Directory service: Automatically configuring.....	154
Creating an Active Directory service: Manually configuring.....	157
Troubleshooting AD connections.....	158
Connecting to an LDAP service	161
LDAP connection settings.....	162
Creating an LDAP Directory service: Automatically configuring.....	164

Creating an LDAP Directory Service: Manually configuring.....	167
Creating an MDM directory service.....	170
Creating a Token authentication service	170
Creating a Kerberos Authentication service	170
Setting up MSCHAPv2 authentication on LDAP	170
MSCHAPv2 termination using an LDAP Password Attribute	171
Creating an NT-Hashed password to support MSCHAPv2 termination against LDAP	173
MSCHAPv2 termination using a Novell universal password	174
MSCHAPv2 termination using an OID authentication descriptor.....	175
Creating a Radius Proxy Authentication service	176
Managing Directory services.....	176
Editing Directory Service configurations.....	176
Renaming a Directory Service.....	177
Deleting a Directory Service.....	179
Directory Sets.....	179
Directory sets panel.....	180
Adding a directory set	181
Renaming a Directory set.....	182
Deleting a Directory set.....	183
Adding directories and authentication servers to a directory set.....	183
Setting Fallthrough rules to handle lookup and authentication failures.....	185
Troubleshooting user lookup and authentication	186
Checking directory service connections.....	186
Checking the Group Cache.....	188
Testing a Directory Service connection.....	188
Advanced troubleshooting for Directory Services and Sets	189
Chapter 11: Mobile Device Management.....	193
Connecting to an MDM service.....	193
Creating an MDM directory service.....	193
Checking an MDM directory service.....	196
MDM enrolled devices.....	196
MDM Enrolled Devices panel.....	196
Filtering the MDM enrolled devices list.....	197
Viewing an MDM enrolled device.....	198
Exporting MDM enrolled device records.....	199
MDM access policies.....	199
Chapter 12: Authentication service.....	201
Setting up a Kerberos Authentication Service	201
Overview of Token Authentication in Ignition Server.....	202
How a user logs in with Token Authentication	202
Components required for Token Authentication	203
Configuring token authentication in Ignition Server.....	203
Prepare your Token Server	203

Connect Ignition Server to RSA Authentication Manager.....	204
Connect Ignition Server to another type of Token Server.....	205
Add the Authentication Server to your Directory Set.....	206
Set Up an Access Policy that uses Token Authentication	207
Prepare Your Authenticators.....	209
Connect Ignition Server to your Authenticators.....	209
Make sure token-capable clients are installed.....	210
Perform an End-to-End test	210
Setting up a RADIUS proxy server.....	211
Use case examples.....	212
Creating a Directory Set.....	214
Adding the RADIUS Proxy Server to a Directory Set.....	214
Creating a RADIUS Access Policy for RADIUS Proxy Server.....	215
Creating a new RADIUS Proxy Policy	215
Creating a RADIUS proxy authentication service.....	217
Configuring the remote RADIUS server	218
Proxying of MAC authentication requests.....	219
Editing Authentication Service Configurations	219
Renaming an Authentication Service.....	219
Deleting an Authentication Service.....	220
Managing a SecurID Authentication Service.....	220
Handling changes to the Node Secret.....	220
Setting Up Supplicants and Authenticators for SecurID Authentication	222
Chapter 13: Virtual Groups and Attributes	223
Introduction to Virtual Groups and Attributes.....	223
How Ignition Server handles multiple Directories	223
Virtual Groups	224
Browsing Virtual Groups.....	224
Mappable Group types for Ignition Server Virtual Groups.....	225
Adding a new Virtual Group	226
Mapping Groups from a Directory Service.....	227
Renaming a Virtual Group.....	228
Deleting a Virtual Group.....	229
User Virtual Attributes	229
About User Virtual Attributes.....	229
Browsing User Virtual Attributes.....	230
Adding a new User Virtual Attribute	231
Mapping Directory Service Attributes to User Virtual Attributes	231
Renaming a User Virtual Attribute.....	233
Deleting a User Virtual Attribute.....	234
Device Virtual Attributes	234
Browsing Virtual Attributes for Devices	234
Adding Virtual Attributes for Devices	235

Chapter 14: User authentication policy	237
Additional policy topics.....	237
How Ignition Server authenticates and authorizes a user	237
Introduction to Policy Management	238
What happened to service categories?	238
Access Policy panel.....	239
Understanding authentication policy	240
One policy allows many authentication protocols	240
Supported authentication types.....	240
Authentication Policy window.....	242
EAP-TLS authentication	243
Factors that limit your choice of a Protocol Credential Certificate	244
Creating an authentication policy	244
Understanding Identity Routing Policy	246
How Ignition Server looks up a user for Authentication and Authorization.....	246
Creating an Identity Routing policy	248
Additional notes on Realm-Matching rules	251
Chapter 15: User Authorization Policy	252
Introduction to User Authorization Policies	252
Structure of user Authorization Policies	252
Elements of a User Authorization Policy	253
How Ignition Server evaluates a user Authorization Policy	253
Reading the rule summary.....	255
Attributes used in Rule Constraints	256
Using Time and Date in a rule.....	260
Conjunctions used to assemble constraints Into a rule	262
Comparison operators for rules.....	263
Creating a RADIUS user authorization policy	264
Enabling or disabling rules within a policy	269
Copying an authorization rule.....	270
Creating an authentication-only policy.....	271
Modifying the default rule to make It authentication-only.....	272
Using a device attribute in a rule	272
Chapter 16: Provisioning policy	274
Introduction to Session Provisioning	274
Setting up session provisioning	275
Vendors panel.....	275
Outbound Attributes	276
Finding a global outbound attribute.....	276
Creating a global outbound attribute.....	276
Overriding the outbound attribute type for one or more authenticators.....	278
Outbound value.....	279
Finding an outbound value.....	279

Creating an outbound value.....	279
Built-in outbound values.....	281
Setting a Provisioning Value.....	282
Inbound Attributes.....	286
Preparing an inbound Attribute for use in an Authorization Rule.....	286
Finding an Inbound Attribute.....	287
Creating a Global Inbound Attribute.....	288
Creating a Vendor-Specific Inbound Attribute.....	289
Device Templates.....	291
Device template window.....	291
Finding a device template.....	292
Creating a Device Template.....	292
Modifying a Device Template.....	294
Applying a device template to your authenticator.....	294
Listing Ignition Server’s set of available RADIUS Attributes	295
Adding a new RADIUS Attribute.....	296
Listing Ignition Server’s set of available VSA attributes.....	297
Adding new VSA.....	298
Adding equipment vendor.....	298
Provisioning FAQ.....	300
Chapter 17: Client posture policy.....	301
How Ignition Server checks client posture.....	301
Enabling NAP on a Windows machine	302
Enable NAP services on the client	302
Enable enforcement on the client.....	302
Configure authentication methods.....	303
Configuring NAP posture profiles.....	304
Chapter 18: VLAN Assignment.....	307
Creating a policy that assigns users to VLANs	307
Create the Outbound Attribute	307
Chapter 19: Windows Machine authentication.....	316
Introduction to Windows Machine authentication	316
Supported authentication protocols	316
Session behavior for Windows Machine authentication	317
NAP support for Peap.....	317
Setting up Microsoft Windows Machine Authentication.....	318
Machine authentication based on ObjectClass	319
Set up Ignition Server to retrieve the objectClass value	319
Write your policy rule	320
Add user policies.....	321
Set up your supplicants	322
Machine authentication based on OU, O, or group membership	322
Prepare your entries in AD	322

Set your user root DN.....	323
Set Ignition Server to retrieve the group membership information	323
Write your policy rule	323
Add user policies.....	325
Set up your supplicants	325
Setting TTL for Windows Machine authentication	325
Chapter 20: TACACS+ authorization.....	327
Introduction to TACACS+ Access Control.....	327
Privilege-level TACACS+ authorization.....	327
Per-command TACACS+ authorization.....	328
Getting started.....	328
Installing your TACACS+ license.....	329
Turning on the Ignition Server TACACS+ service.....	329
Creating a Command Set.....	331
Viewing or editing a command set.....	332
Creating a TACACS+ Access Policy.....	333
Enable your devices for TACACS+ authorization.....	335
Using the TACACS+ global authenticator.....	337
Chapter 21: MAC Authentication.....	339
Introduction to MAC Authentication.....	339
Creating a MAC-Auth policy.....	340
Setting Up MAC Auth.....	341
MAC authentication set-up procedure example.....	343
VLAN assignment using the Device Record VLAN fields.....	348
Allowed MAC Address formats.....	351
Notes on writing MAC authorization rules.....	351
Comparing a Device's MAC Address.....	352
Checking a Device's Group Membership.....	352
Chapter 22: Asset Correlation.....	353
Introduction to Asset Correlation.....	353
MAC Address vs. Windows Machine Authentication	353
Creating Asset Correlation policies.....	354
Requiring the user to connect using an Allowed Device.....	355
Requiring the user to connect using his or her Assigned Device.....	356
Requiring the user to connect using a Machine Authenticated-Device.....	358
Viewing currently Authenticated Devices	360
Chapter 23: Command Line Interface.....	361
Connecting to the CLI through an SSH connection	361
Connecting via SSH.....	362
Appendix A: Installing Ignition Server.....	364
Installation prerequisites.....	364
Map out your Ignition Server Deployment	364
VMware ESXi server.....	365

Preventing automatic VMware tools updates.....	366
Checking the VMware Tools status (ESXi 5.1 and up).....	367
Importing VM.....	368
Applying the license.....	376
Installing the license.....	377
Installing the Ignition Dashboard desktop application.....	378
Connect Ignition Server for the first time.....	388
Run Ignition Dashboard.....	388
Install Your Ignition Server Licenses.....	389
Configure the Ignition Server.....	389
Uninstalling Ignition Dashboard.....	390
Appendix B: Paired server high availability configuration.....	391
HA terminology	391
Overview of HA Pairs.....	391
Creating an HA Pair	392
Start and connect the Ignition Server	392
Run the HA Wizard.....	393
Bind Services to the VIP.....	401
Restoring a saved VIP configuration	402
Managing an HA Pair	403
Managing Virtual Interfaces (VIPs).....	405
Breaking an HA pair using Dashboard.....	408
Breaking an HA Pair using the CLI.....	409
Reconnecting a Broken HA Pair.....	409
Reinitializing Nodes in an HA Pair	409
Backing Up Data on an HA Pair	410
Restoring Data on an HA Pair	410
Updating Firmware on an HA Pair	411
Replacing an Ignition Server in an HA Pair	411
Changing the IP Address of the Admin Port or HA Port in an HA Pair	413
Restoring a Non-Responsive Unit in an HA Pair	414
Appendix C: Extended High Availability configuration.....	416
Overview.....	416
Configuring scheduled exports.....	417
Configuring scheduled imports.....	419
Editing an export or import schedule.....	421
Appendix D: Backup and Restore Procedures.....	422
Introduction to Ignition Server Backup and Restore	422
Creating a backup.....	422
Troubleshooting.....	423
Configuring scheduled backups.....	423
Restoring from a backup file.....	425
Admin Port IP Address	426

Restoring data from a backup file.....	426
Troubleshooting.....	427
Appendix E: Firmware Update Procedures.....	429
Checking the Firmware version	429
View the current Firmware version	429
View the current Firmware version and all installed Images and Packages.....	430
Loading a Firmware Image or Package	431
Activating a firmware image or package	432
Activating a Firmware Image on an HA Pair	434
Deleting a Firmware file	434
Viewing Image and Package information	435
Upgrading Ignition Server.....	436
Upgrading Ignition Server	437
Appendix F: Setting up logging.....	438
Setting User Preferences.....	438
Setting Up Ignition Server Logging.....	438
Turning on logging.....	438
Setting the Level of Logging to be recorded	440
Setting up FTP log export	440
Exporting Logs.....	441
Directing log messages to a Syslog server	443
Sending log messages Via E-Mail.....	444
Monitoring Ignition Server via SNMP.....	445
Configuring Ignition's SNMP Service	445
Connecting to Ignition's SNMP Service.....	447
Example SNMP Queries.....	447
Data Objects exposed by the Ignition Server SNMP Service	448
Port names used in SNMP output.....	449
Appendix G: Viewing logs and statistics.....	450
Overview of Logging and Log types	450
Viewing and managing logs.....	451
Viewing Logs.....	451
Filtering your view of the Logs	451
Criteria for filtering Log messages	452
Managing the stored logs	453
Log Viewer.....	453
Access Log: RADIUS and TACACS+ Accounting.....	454
Access Record Details.....	456
Audit Log.....	460
Security Log.....	462
System Log.....	463
Statistics tab	464
Transactions tab.....	465

Directory Services tab.....	466
Protocols tab.....	467
System Health tab.....	468
Contents of the System Health Tab.....	468
Viewing the System Health Tab.....	468
Directory Services Status Tab.....	468
AAA Summary tabs.....	469
Contents of the AAA Summary Tabs.....	469
Viewing the AAA Summary Tabs.....	470
Specifying the number of records to be shown.....	471
User Accounting tab.....	471
Contents of the User Accounting tab.....	471
Viewing the User Accounting Tab.....	473
Learned Devices tab.....	473
Contents of the Learned Devices Tab.....	473
Viewing the Learned Devices Tab.....	474
Filtering the Learned Devices Tab.....	474
Revoking the session of a Machine-Authenticated Device	474
Debug Logs.....	474
SAML Access Summary tab.....	475
Contents of the SAML Access Summary tab.....	475
Viewing the SAML Access Summary tab.....	475
Administration Sessions tab	476
Administration Access Summary tab.....	476
Appendix H: Troubleshooting	478
Generating a trouble ticket	478
Troubleshooting common problems.....	479
Problem: Cannot connect to Ignition Dashboard.....	479
Problem: Connecting Dashboard to Ignition Server Fails.....	480
Problem: Ignition Server fails to respond to RADIUS and/or TACACS+ requests.....	482
Problem: Authorization policy stops working unexpectedly.....	482
Problem: Authentication fails on Active Directory.....	484
Problem: HA Set-up fails.....	484
Problem: Two primary HA nodes detected.....	485
Problem: Errors occur during Directory Service Set-Up	486
Problem: Unable to Map Virtual Attributes.....	487
Problem RADIUS Proxy Service Fails.....	487
Glossary	489

Chapter 1: Introduction

Purpose

Avaya Identity Engines Ignition Server Administration explains how to configure 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 configuring 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.

Related resources

Documentation

See the following related documents.

Title	Purpose	Document number
<i>Avaya Identity Engines Ignition Server Getting Started</i>	Installation and simple configuration	NN47280–300
<i>Configuring and Managing Avaya Identity Engines Single-Sign-On</i>	Configuration, management, and deployment	NN47280–502
<i>Avaya Identity Engines Ignition Guest Manager Configuration</i>	Installation, configuration, and management	NN47280–501
<i>Avaya Identity Engines Ignition CASE Administration</i>	Installation, configuration, and deployment	NN47280–603
<i>Avaya Identity Engines Ignition Access Portal Administration</i>	Installation, configuration, and deployment	NN47280–604
<i>Avaya Identity Engines Ignition Analytics</i>	Installation, configuration, and maintenance	NN47280–601
<i>Avaya Identity Engines Ignition Server Release Notes</i>	Reference	NN47280–400

Training

Ongoing product training is available. For more information or to register, you can access the Web site at <http://avaya-learning.com/>.

Viewing Avaya Mentor videos

Avaya Mentor videos provide technical content on how to install, configure, and troubleshoot Avaya products.

About this task

Videos are available on the Avaya Support website, listed under the video document type, and on the Avaya-run channel on YouTube.

Procedure

- To find videos on the Avaya Support website, go to <http://support.avaya.com> and perform one of the following actions:
 - In **Search**, type `Avaya Mentor Videos` to see a list of the available videos.
 - In **Search**, type the product name. On the Search Results page, select **Video** in the **Content Type** column on the left.
- To find the Avaya Mentor videos on YouTube, go to www.youtube.com/AvayaMentor and perform one of the following actions:
 - Enter a key word or key words in the Search Channel to search for a specific product or topic.
 - Scroll down Playlists, and click the name of a topic to see the available list of videos posted on the website.

 **Note:**

Videos are not available for all products.

Subscribing to e-notifications

Subscribe to e-notifications to receive an email notification when documents are added to or changed on the Avaya Support website.

About this task


You can subscribe to different types of general notifications, for example, Product Correction Notices (PCN), which apply to any product or a specific product. You can also subscribe to specific types of documentation for a specific product, for example, Application & Technical Notes for Ethernet Routing Switch 8800.

Procedure

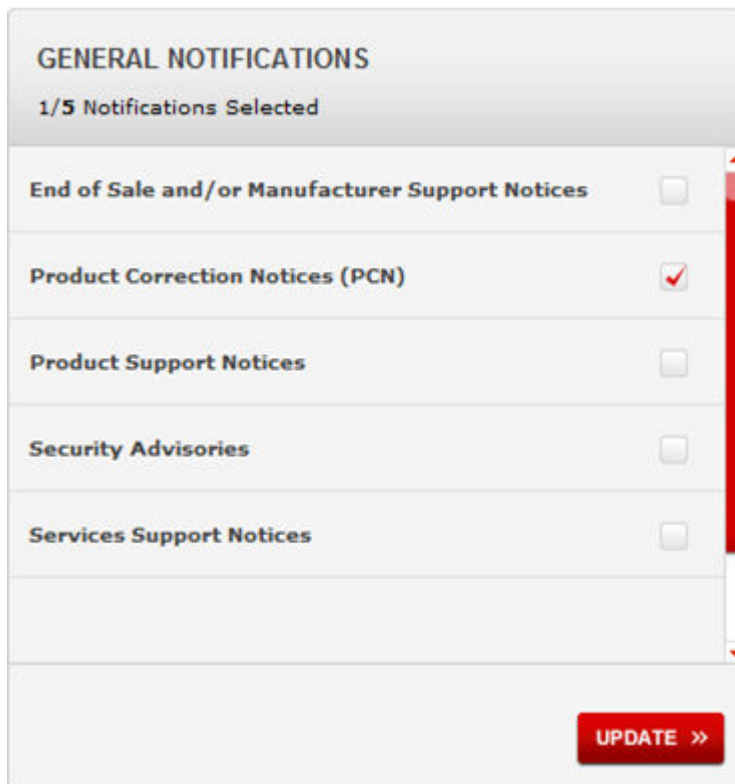
1. In an Internet browser, go to <https://support.avaya.com>.
2. Type your username and password, and then click **Login**.
3. Click **MY PROFILE**.




4. On the site toolbar, click your name, and then click **E Notifications**.



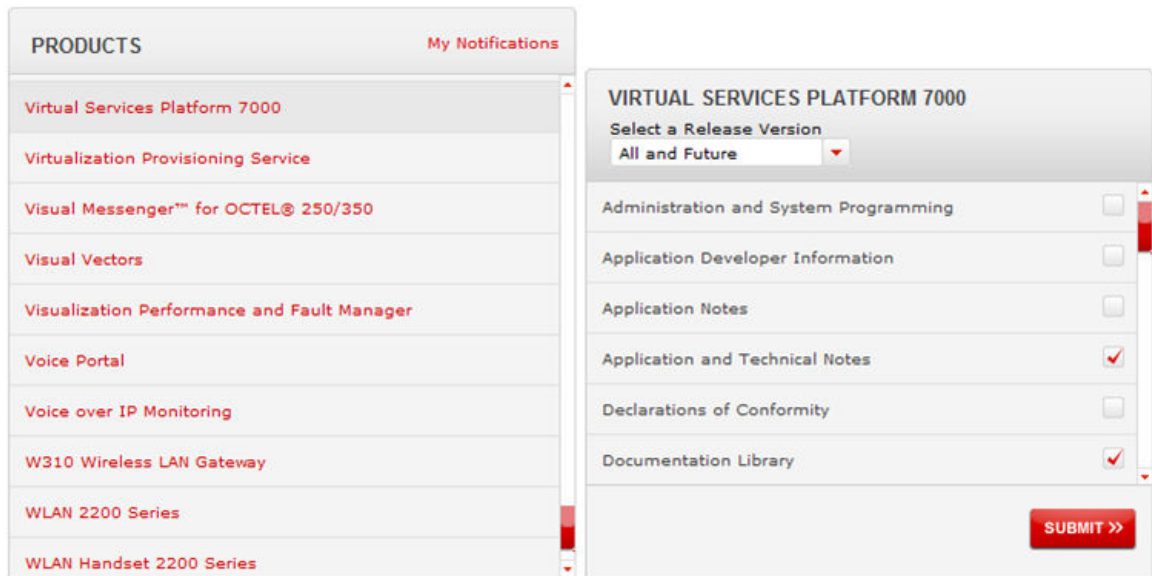
5. In the GENERAL NOTIFICATIONS area, select the required documentation types, and then click **UPDATE**.



6. Click **OK**.
7. In the PRODUCT NOTIFICATIONS area, click **Add More Products**.



8. Scroll through the list, and then select the product name.
9. Select a release version.
10. Select the check box next to the required documentation types.



11. Click **Submit**.

Searching a documentation collection

On the Avaya Support website, you can download the documentation library for a specific product and software release to perform searches across an entire document collection. For example, you can perform a single, simultaneous search across the collection to quickly find all occurrences of a particular feature. Use this procedure to perform an index search of your documentation collection.

Before you begin

- Download the documentation collection zip file to your local computer.
- You must have Adobe Acrobat or Adobe Reader installed on your computer.

Procedure

1. Extract the document collection zip file into a folder.
2. Navigate to the folder that contains the extracted files and open the file named `<product_name_release>.pdx`.

3. In the Search dialog box, select the option **In the index named <product_name_release>.pdx**.
4. Enter a search word or phrase.
5. Select any of the following to narrow your search:
 - Whole Words Only
 - Case-Sensitive
 - Include Bookmarks
 - Include Comments
6. Click **Search**.

The search results show the number of documents and instances found. You can sort the search results by Relevance Ranking, Date Modified, Filename, or Location. The default is Relevance Ranking.

Support

Go to the Avaya Support website at <http://support.avaya.com> for the most up-to-date documentation, product notices, and knowledge articles. You can also search for release notes, downloads, and resolutions to issues. Use the online service request system to create a service request. Chat with live agents to get answers to questions, or request an agent to connect you to a support team if an issue requires additional expertise.

Chapter 2: New in this release

The following sections detail what is new in *Avaya Identity Engines Ignition Server Administration* for Release 9.1.

Features

See the following sections for information about feature changes.

Mobile Device Management support

Identity Engines Release 9.1 adds support for Mobile Device Management (MDM) by integrating third-party MDM services. Identity Engines Release 9.1 supports Citrix XenMobile MDM.

The MDM feature provides more control for and secure access to bring your own device (BYOD) deployments in a corporate network. With the new MDM feature, mobile devices, whether they are corporate-owned or personal, are enrolled in an MDM server. The device attributes such as the OS version are stored in the MDM server and indexed by the MAC address of the device.

The Ignition Server interfaces to the different MDM services to collect the device attributes and save them in the Internal Store. Device attribute lookup happens locally on the Ignition Server. During user authentication, the device attributes are evaluated and fed to the policy engine in the final AAA authorization decision making.

Avaya Identity Engines Ignition Access Portal enhancements

Avaya Identity Engines Ignition Access Portal Release 9.1 contains numerous updates including multi-IN interfaces, multi-OUT interfaces, multi-Success Pages, multi-Captive Portal Zones, Access Groups, and enhanced device fingerprinting, which you can limit by a time lapse period from initial device fingerprint through Avaya Identity Engine Ignition Server. For more information, see *Administering Avaya Identity Engines Ignition Access Portal*, NN47280–604.

Extended High Availability

Release 9.1 introduces the first phase of an Extended High Availability (HA) configuration with geographically redundant Identity Engines (IDE) Ignition servers. One site is designated as a root site, which contains the primary, active Ignition Server. One or more sites are designated as branch sites, which contain secondary, inactive Ignition Servers. Configuration of user devices and guest user accounts occurs on the root site, and periodic exports of this data transfers the user and device information to the branch site. In the event of root site failure, a branch site can be made active to process access requests.

Guest Manager enhancements

Guest Manager Release 9.1 includes the following:

- A new sponsor approval workflow for self-service guest users. Administrators configure whether or not a sponsor's approval is required. If required, the guest user creates the self-service account and the account is approved by the sponsor before access is granted to the guest user.
- Enhancements to the self-service template to make it simpler with fewer fields and more automation.
- Enhancements to the email SMTP and SMS gateway configurations.
- A Trouble Ticket option on the Administrator GUI that generates a trouble ticket file to be sent to Avaya technical support in the event of a Guest Manager fault.
- Guest Manager is now available as a virtual appliance with OS RHEL instead of a Windows application.

Advanced RADIUS proxy

Ignition Server Release 9.1 introduces Advanced RADIUS Proxy. Advanced RADIUS Proxy enhances the Identity Engines RADIUS Proxy feature to allow forwarding RADIUS authentication requests to a remote RADIUS server and manipulating the response by adding, deleting, or modifying the RADIUS attributes returned by the remote RADIUS server before delivering them to the Authenticator that originated the RADIUS authentication request.

A key use case of the Identity Engines Advanced RADIUS Proxy is EDUROAM (Education Roaming). EDUROAM is a secure, world-wide roaming access service that allows students and staffs from participating institutions to obtain network access in the campus when they visit other participating institutions. The core idea is that the authentication server at the new campus would simply proxy the authentication requests to the other authentication server where the user account is originally provisioned. Identity Engines Ignition Server Intelligent RADIUS Proxy feature is enhanced to act as a RADIUS proxy server and forwards the authentication requests to a remote RADIUS server to support EDUROAM while providing the administrator full control of network access.

Another key use case of the Identity Engines Advanced RADIUS Proxy is to allow overlaying the Ignition Server over any network infrastructure and allow forwarding authentication requests to an existing NAC system. The customer may choose not to replace the NAC system, or may choose to gradually migrate from their existing NAC system to Identity Engines over time.

The new Advanced RADIUS Proxy:

- Introduces a formal terminology of Forwarding server and Remote server. A Remote server can be any server that is capable of handling RADIUS.
- Supports advanced capability such as delete/update/append (to) attributes sent by the Remote server for local authorization enforcement based on local conditions.
- Supports hierarchical roaming server deployment as recommended by EDUROAM.
- Provides access and accounting logs on Remote server activities, on the Forwarding server.
- Provides the ability to proxy based on exact or partial realm matches (for example, user@sv.avaya.com or user@avaya.com).
- Allows override of Remote server reject responses and allows controlled access locally if required.
- Supports advanced Guest Roaming capability, allowing dynamic realm-based authentication.

Chapter 3: Introduction to Avaya Identity Engines Ignition Server

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

Installation prerequisites, installation procedures, and the initial connection to the Avaya Identity Engines Ignition Server from Avaya Identity Engines Ignition Dashboard are described in [Installing Ignition Server](#) on page 364.

What is Avaya Identity Engines Ignition Server?

Avaya Identity Engines Ignition Server is an enterprise grade network access policy server. The Ignition server is also an 802.1X-capable RADIUS authentication server that grants users and their devices different access levels, or denies users access to your network based on your access policies. Use the Ignition Server to create a single set of policies that control access for all of the ways that users connect: through wired, wireless, or VPN. Ignition Server stores access policies, while user accounts remain in your traditional user store(s), such as such as Microsoft Active Directory, Open LDAP, Novell eDirectory, RSA Authentication Server, and others.

Ignition Server includes an easy-to-configure policy engine that lets you make network access decisions based on the user's identity, account details and group memberships, location of the login attempt, time of day, and other pieces of information. For example, an Ignition Server policy can grant users access based on their identity, their point of access (which network switch or WAP they are connecting through), and their laptop security state (ensuring their laptop is a company-owned laptop as recorded in the corporate Active Directory store).

Ignition Server's abilities to check whether a user's workstation has passed MAC authentication and Windows machine authentication are key features that set it apart from other network access control tools. Ignition Server lets you combine many policy elements to enforce a single rule, such as how to authenticate a user with PEAP/MSCHAPv2, check that their device has been authenticated, and if those are successful, assign the user to the appropriate VLAN based on their role. Ignition Server also authenticates devices. You can configure Ignition Server to offer a bypass of 802.1X authentication for older devices on your network that cannot perform an 802.1X authentication by using the Ignition Access Portal.

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:** 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 configured 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 can deploy two Ignition Servers as a linked pair that offers a highly available RADIUS service. You can also exchange user and device details between geographically dispersed Ignition Servers for Extended high availability.

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, and so on.). 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, and a geographically redundant 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 Authentication, Authorization, and Accounting (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 can include passwords, 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.

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.

Table continues...

Feature	Ignition Server Action
Network and port segmentation	<p data-bbox="451 239 1451 394">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 62. 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.</p> <p data-bbox="451 415 1406 474">As shown in Configuring the Ignition Server's network ports on page 71 you may place limits on what traffic each Ignition Server ports may carry.</p>

Chapter 4: 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 [Installing Ignition Server](#) on page 364.

For information on additional Ignition Server management operations, see the following:

- [Backup and Restore Procedures](#) on page 422
- [Firmware Update Procedures](#) on page 429
- [Setting up logging](#) on page 438

Administration roles

Identity Engines supports multiple administration roles (other than the System Administrator). These roles allow the system administrator to define different permission masks for different users on the system. The users can only be assigned to the group by the System Administrator, and can be either an internal user (existing in the local store) or an external user (from a Directory Service or Directory Set). Users can belong to only one group.

The groups are as follows:

- Configuration Administrator
- Troubleshooting Administrator
- Monitoring Administrator

For information on how these administration roles are managed in Ignition Dashboard, see [Administration functionality](#) on page 42.

System Administrator

The System Administrator has full access to Dashboard.

There can only be one System Administrator.

Configuration Administrator

The Configuration Administrator has full access to most of the items under Site Configuration. Site, node, license, certificate, server, and RBAC management-related functions are not accessible to the Configuration Administrator. Those functions can only be performed by the System Administrator. There can be multiple users assigned to the Configuration Administrator group; however, only one Configuration Administrator can be logged in at a time.

The Configuration Administrator has all the system permissions that the Troubleshooting Administrator and Monitoring Administrator have.

Troubleshooting Administrator

The Troubleshooting Administrator has full access to the Dashboard Monitor and is able to browse the Dashboard Configuration. They can also access the Dashboard Troubleshoot functions. Multiple Troubleshooting Administrators can login simultaneously without impacting another user's login.

The Troubleshooting Administrator has all the system permissions that the Monitoring Administrator has.

Monitoring Administrator

The Monitoring Operator has full access to the Dashboard Monitor and is able to browse the Dashboard Configuration, but cannot make any permanent or temporary configuration changes that impact the network access behavior of Identity Engines. Multiple Monitoring Administrators can login simultaneously without impacting another user's login.

Monitoring administrators cannot use the troubleshooting functionality.


Launching Ignition Dashboard

Run Ignition Dashboard.

Procedure

1. Double-click the Ignition Dashboard icon on your desktop or select **Start > Programs > Ignition Dashboard > Ignition > Dashboard**.

The following login window appears.



2. Type the system 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 System Administrator login name](#) on page 56.
3. In the **Connect To** field, do one of the following:
 - a. To connect to an individual Ignition Server site, type the hostname or IP address of your Ignition Server.
 - b. 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 37.
 - c. Click **OK**. If you are unable to log in, see [Problem: Cannot connect to Ignition Dashboard](#) on page 479.

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 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

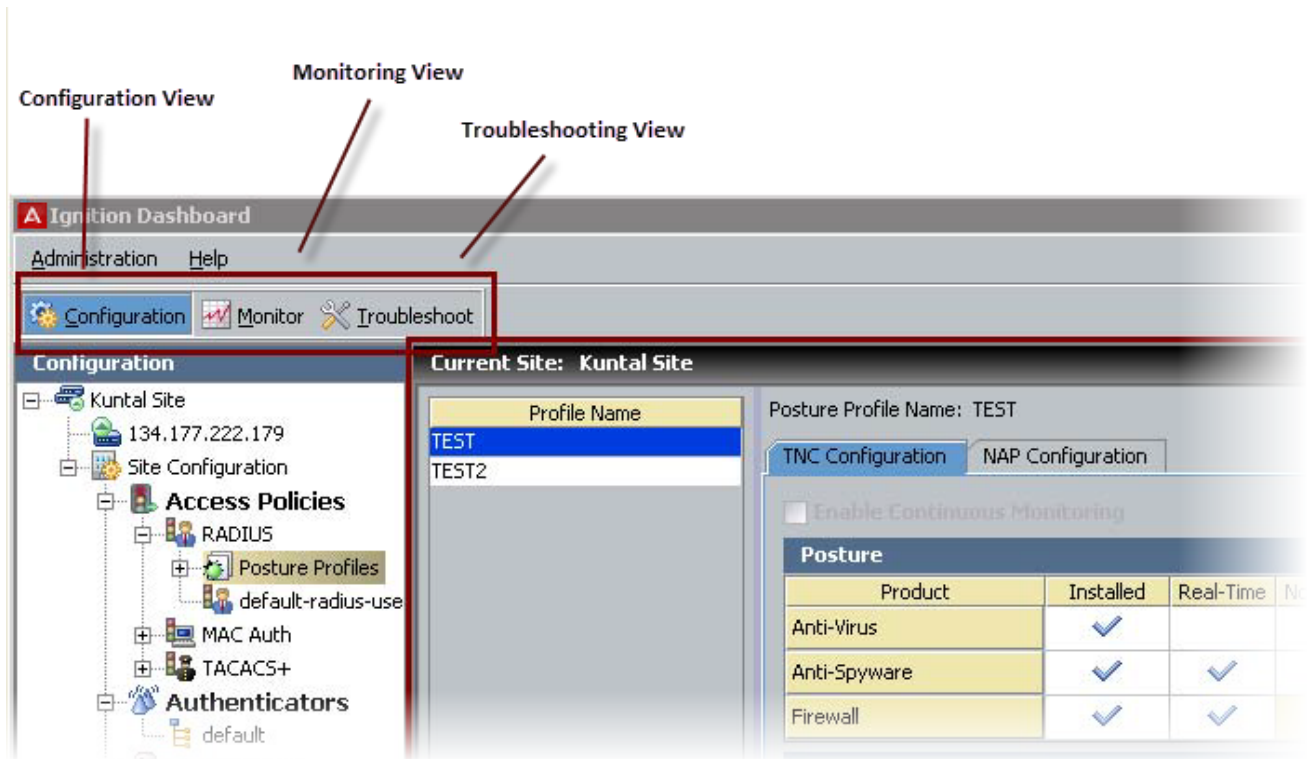
If an error occurs on the Ignition Server, Dashboard displays an alert icon at the bottom of the Dashboard main window, as shown below. Click the alert icon to view a dialog window showing a description of the error. Before you dismiss the dialog, you should clear the message by clicking on the message 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 following figure illustrates Ignition Dashboard, your graphical user interface for configuring your Ignition Server(s).



After you have logged in to Dashboard, your User name and the Role associated with that user name are displayed in the bottom right corner of the Dashboard status bar.

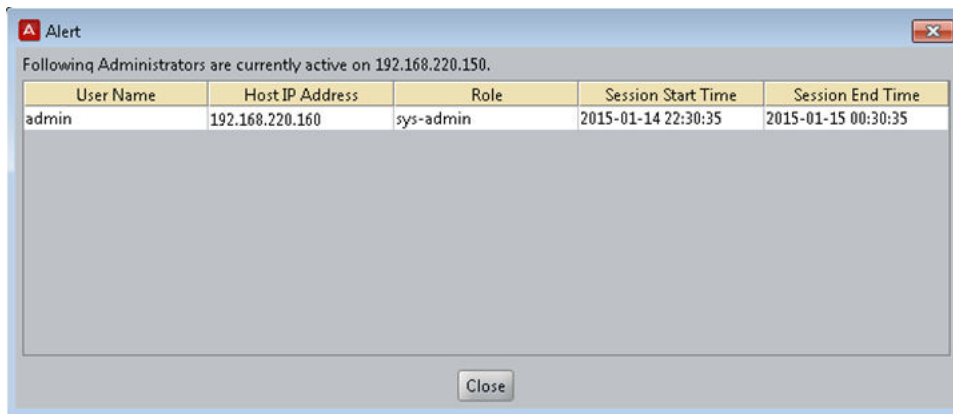
For example, in the following figure, the user name is “mon” and the role associated with the user is “Monitoring Administrator”.



Multiple logins alert

After you have logged in to an appliance using Dashboard, and another user logs in to that same appliance, then a notification displays in your Dashboard giving information about this new user, as shown in the following figure.

The information displayed includes the User Name, the IP Address of the device from which the user logged in (Host IP Address), the Role of that user and their Session Start Time and Session End Time.



Alert

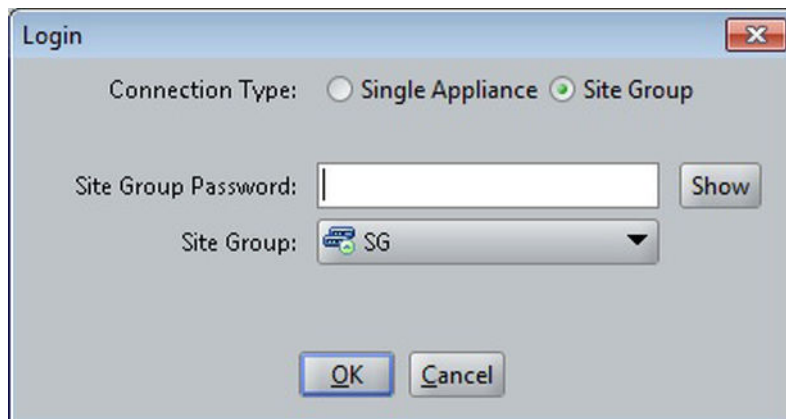
Following Administrators are currently active on 192.168.220.150.

User Name	Host IP Address	Role	Session Start Time	Session End Time
admin	192.168.220.160	sys-admin	2015-01-14 22:30:35	2015-01-15 00:30:35

Close

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 following figure, 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.



Login

Connection Type: Single Appliance Site Group

Site Group Password: Show

Site Group:

OK Cancel

! Important:

Only users with System Administrator credentials can log in to a Site Group.

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.

Procedure

1. Make sure your Ignition Server site has been given a unique name. See [Renaming an Ignition Server site](#) on page 56
2. In the main window of Dashboard (with Dashboard already connected to an Ignition Server), select **Administration > Site Group Management**.

The **Site Group Management** window displays.



3. In the **Site Group Management** window, select **Actions > Add Site Group**.
4. Enter the password to be used when you connect to or configure the site group. To confirm, reenter the password and click **OK**.
5. Type a name for the site group and click **OK**.
6. Add an Ignition Server to the Site Group:
 - a. In the **Site Group Management** window, click the name of your group and select **Actions > Add Site Group**.
 - b. In the **Add Site Group** window, enter the **Site Group Name**. Click **Add** at the bottom of the window.

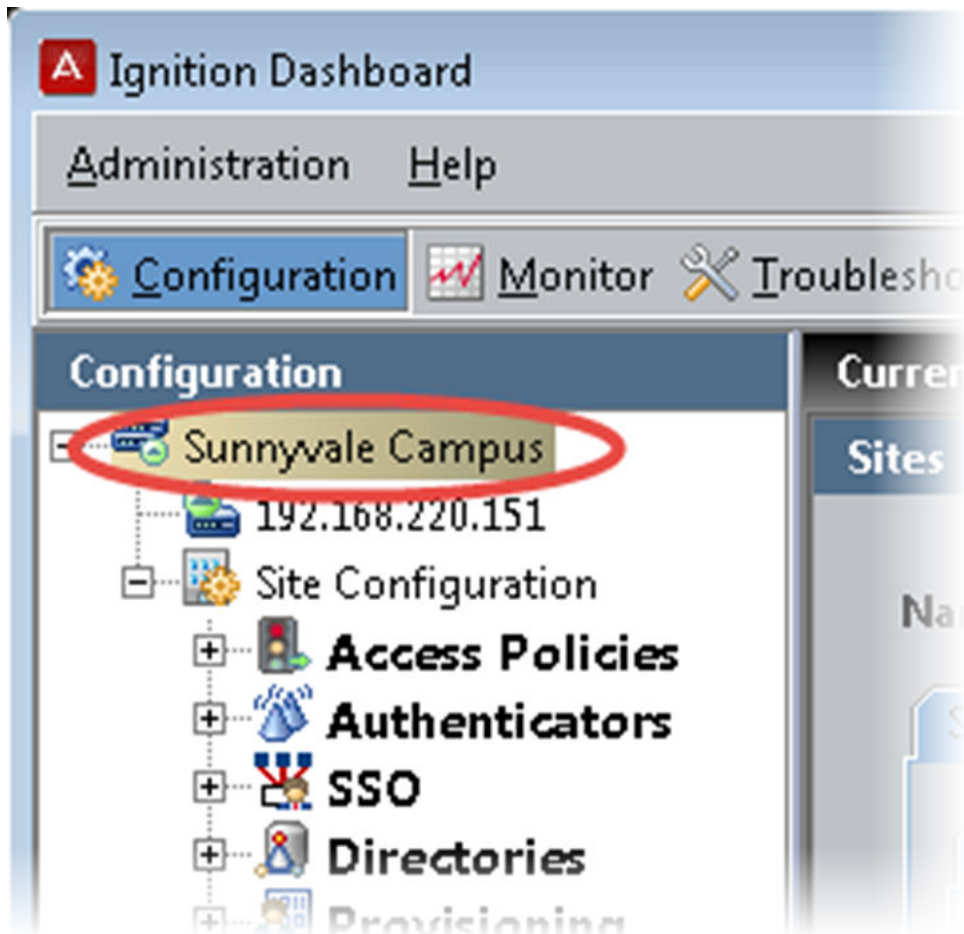
This displays another window where you enter the **IP address, username, and password** for the user you wish to add. Click **OK**.

! Important:

The user you are adding must have System Administrator credentials. Only a System Administrator can log in to the Site Group.

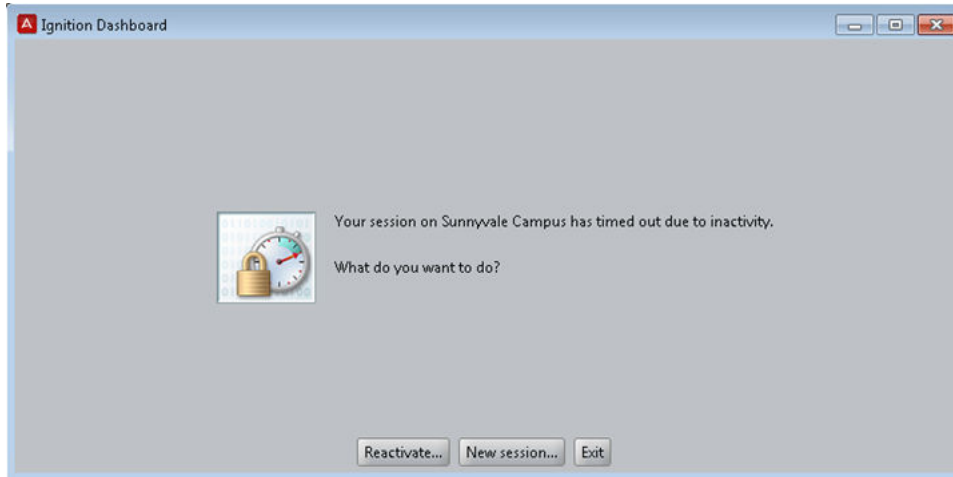
7. Repeat Step 5 for the other appliances in the group.
8. Configure the password for the site group:
 - a. In the **Site Group Management** window, click the name of your group and select **Actions > Modify Configuration Password**.

- b. In the dialog window, type the current password (“admin” is the default) and type the desired new password. Type the new password again to confirm it, and click **OK**.
 9. Click **Close** to close the Site Group Management window.
 10. Disconnect Dashboard from the current appliance and reconnect to the site group:
 - a. In the main window, select **Administration > Logout**.
 - b. Select **Administration > Login**.
 - c. Type the **Password** you configured for the site group in Step 7.
 - d. Click the **Connect To** drop down-box and choose the name for your site group.
 - e. Click **OK** to connect.
 11. The Dashboard Configuration tree lists all the sites in your site group. Click the name of a site to manage that site.



Idle time-out

By default, the Ignition Dashboard user interface is set to time out after a session of Ignition Dashboard has been idle for 20 minutes, unless the **Do not lock Ignition Dashboard** option has been selected in the **Preferences** configuration window. When your current session times out due to inactivity, Ignition Server displays the **Idle Timeout** window.



Choose one of the following actions:

- Click **Reactivate...** to reactivate the current session of Ignition Dashboard on the same Ignition Server. See [Reactivating a session](#) on page 40.
- Click **New Session...** to start a new session. See [Starting a new session](#) on page 41.
- Click **Exit** to exit from the Ignition Dashboard application. See [Quitting from a timed-out session](#) on page 41.

Reactivating a session

In order to reactivate the current session of Ignition Dashboard on the same Ignition Server:

Procedure

1. Click **Reactivate** in the idle time-out window. The **Reactivate** window displays.
2. Enter your administrator password.
3. Click **OK**.

The Ignition Dashboard reappears in the same state as it was before the time-out.


Starting a new session

When you choose 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.

Follow this procedure to start a new session of Ignition Dashboard.

Procedure

1. Click **New Session** in the Idle Timeout window .
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:
 - a. Click **Cancel**. Ignition Server displays the idle time-out window again.
 - b. Click **Reactivate** to reauthorize the current session that timed out.

The Ignition Dashboard reappears in the same state as it was before the time-out.

4. If you want to discard the data associated with the session that has timed out:
 - a. Click **Proceed Without Saving**. Ignition Server drops the connection with the timed-out Ignition Server and displays the **Login** window.
 - b. 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

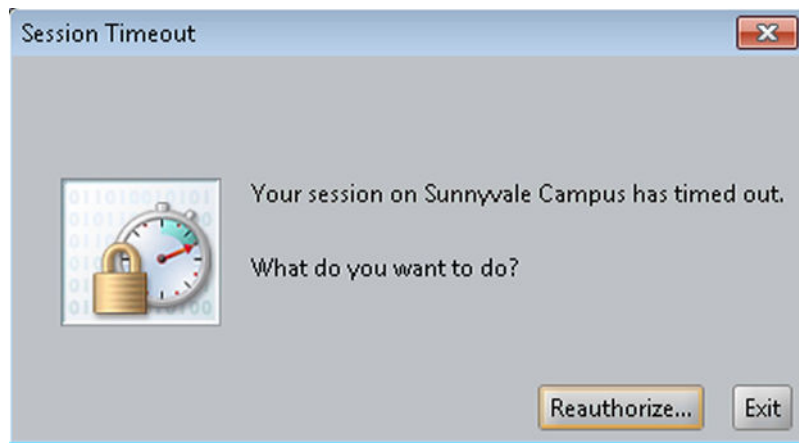
When you click **Exit** in the Idle Timeout window, Ignition Server disconnects the timed-out session of Ignition Dashboard from the Ignition Server, and closes the Ignition Server dashboard application.

Session time-out

Each time a user connects the Dashboard to an appliance, a session is established in the Ignition Server. Session time-out is the duration after which the session established by Dashboard with the Ignition server is terminated. If the Session time-out is 120 minutes, for example, then after two hours of the Dashboard being logged in, the session is now removed from the server. When this occurs, a session time-out message is displayed. This indicates to the user that the initial session established with the appliance is no longer valid.

When the session time-out message displays, choose one of the following options:

- Click **Reauthorize** and establish a new session by entering your credentials.
- Click **Exit**.



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 87.

Administration functionality

The functionality to manage administration roles is available on Ignition Dashboard under **Administration**.

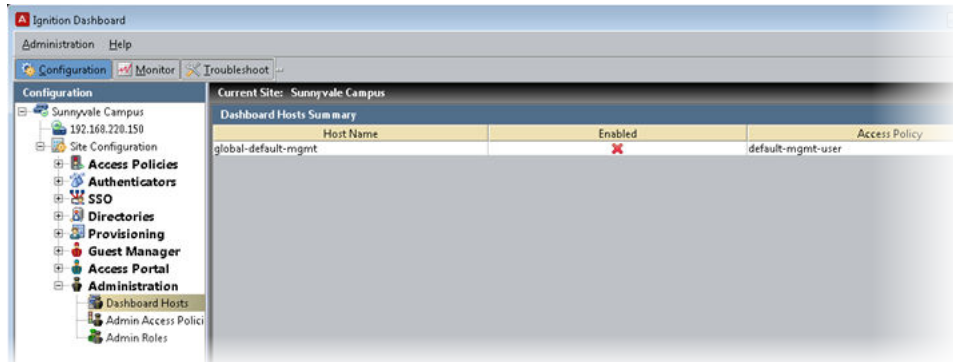
Navigating the Dashboard Hosts

This procedure shows the functionality and information that is available in the Dashboard Hosts windows.

Procedure

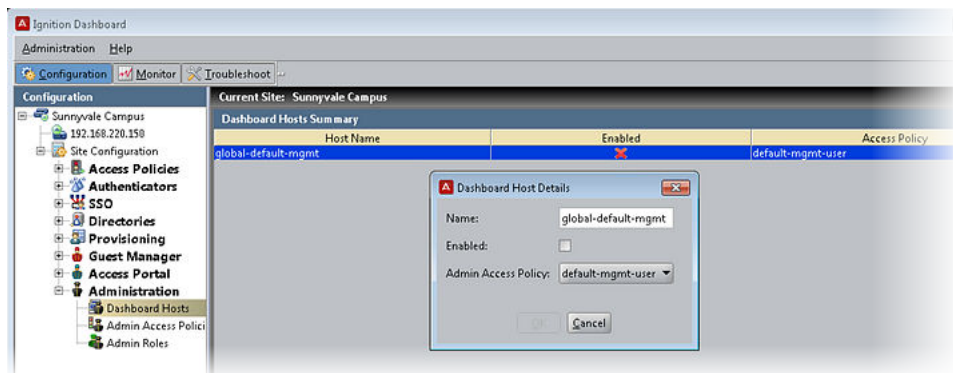
1. In the Configuration View, expand **Administration** and select **Dashboard Hosts**.

The **Dashboard Hosts Summary** panel lists the default management Host.



2. Double-click the dashboard host entry, or highlight the entry and click **Edit**.

The **Dashboard Host Details** dialog box opens, allowing you to edit any of the values assigned to this dashboard host.



3. Make any required changes. Once you begin making changes, an **OK** button appears.
4. Click **OK** when you are finished.

Your changes are displayed on the **Dashboard Hosts Summary** panel.

Navigating the Admin Access Policies

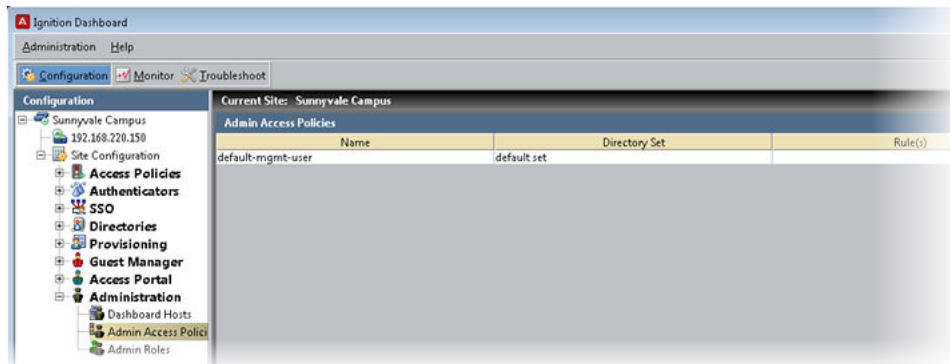
This procedure shows the functionality and information that is available in the Admin Access Policies windows.

Admin Access Policies determine who is granted access and the type of session that is created. The policies are made up of a series of rules that are based on user or system attributes. As an example, the roles could be assigned based on group membership. If a user satisfies the rules of a policy that pertain to a particular level of administrator role, the user attains the level of administrator associated with those rules. The session time-out value and idle time-out value could also be assigned based on the rules.

Procedure


1. In the **Configuration** view, expand **Administration** and select **Admin Access Policies**.

The **Admin Access Policies** panel lists all available policies, including the standard default-mgmt-user policy and any policies you have added.

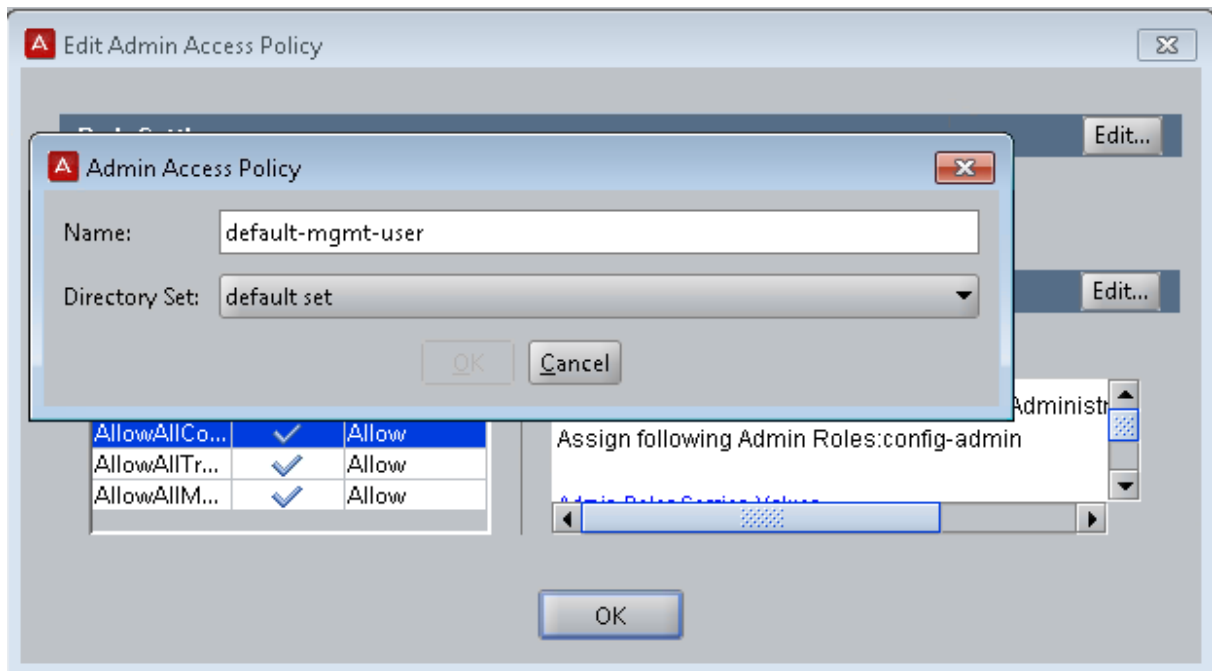


2. Double-click any policy entry, or highlight any entry and click **Edit**.

The **Edit Admin Access Policy** dialog box opens, allowing you to edit any of the values assigned to this policy or to add new rules to the policy.

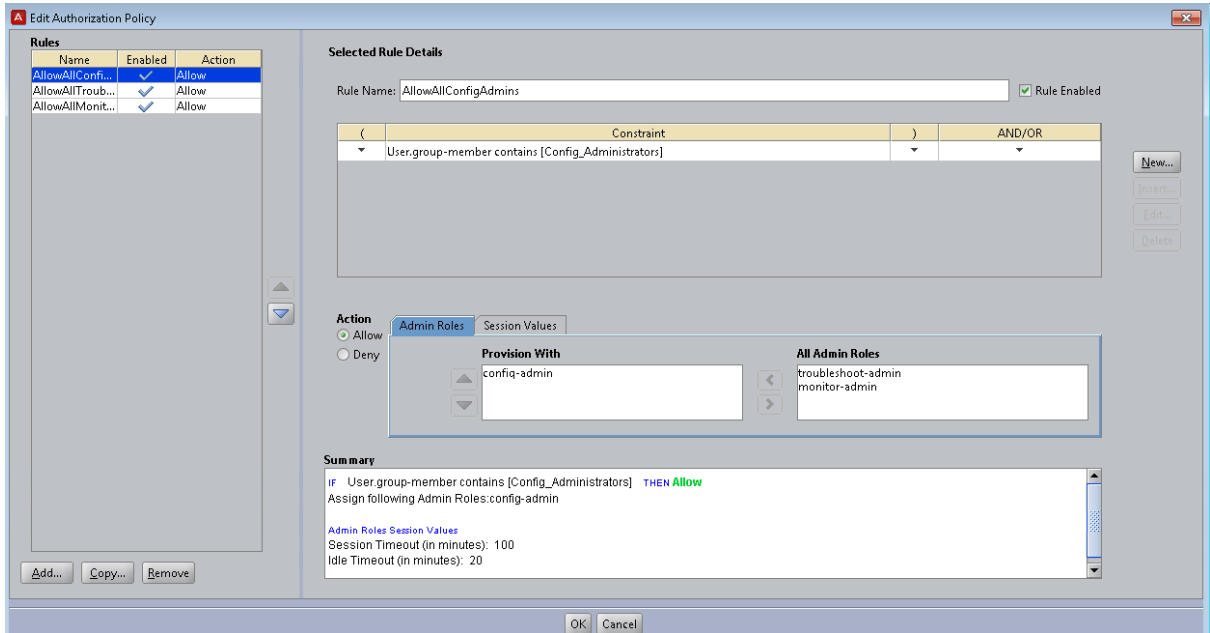


- Click **Edit** on the **Basic Settings** banner to change the policy name or directory set. The **Admin Access Policy** dialog box opens.



- Make any required changes in the Admin Access Policy dialog box. Once you begin making changes, an **OK** button appears. Click **OK** when you are finished.

- Click **Edit** on the **Authorization Policy** banner to change the content of the policy.
The **Edit Authorization Policy** dialog box opens.



The **Rules** panel on the left lists all the individual rules that make up the policy. Using the buttons at the bottom of the panel, you can add a new rule or you can copy or delete an existing rule.

When you highlight a rule in the Rules panel, the **Selected Rules Details** panel on the right displays the details for that rule.

The **Selected Rules Details** panel contains multiple options for editing a policy.

- 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 or the System).
- Each rule has an action to ALLOW or DENY the access request.
- Each rule can have only one Admin Role associated with it. An Admin Role describes what access a user has. For example, a user associated with a config-admin role has permissions to perform the various operations described in [Administration Roles - Configuration Administrator](#) on page 32.

The entire rule is displayed in the **Summary** pane at the bottom.

- Click **OK** at the bottom when you are finished making your changes.

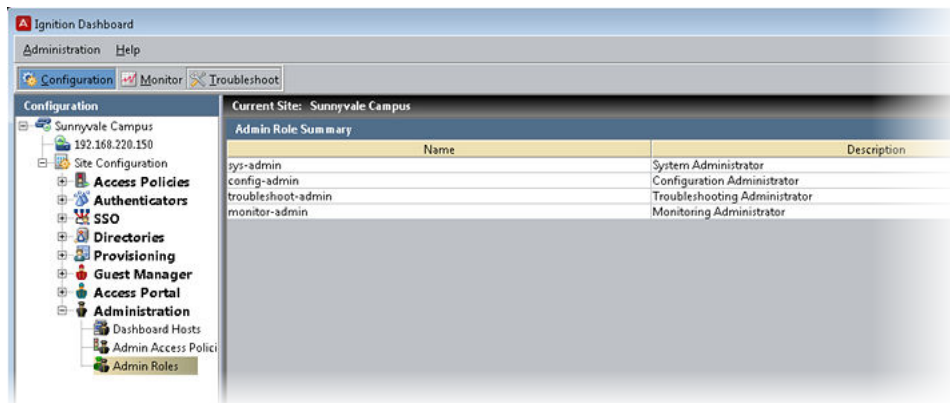
Navigating the Admin Roles

This procedure shows the functionality and information that is available on the Admin Roles windows.

Procedure

1. In the Configuration View, expand **Administration** and select **Admin Roles**.

The **Admin Role Summary** panel lists the existing administrator roles.



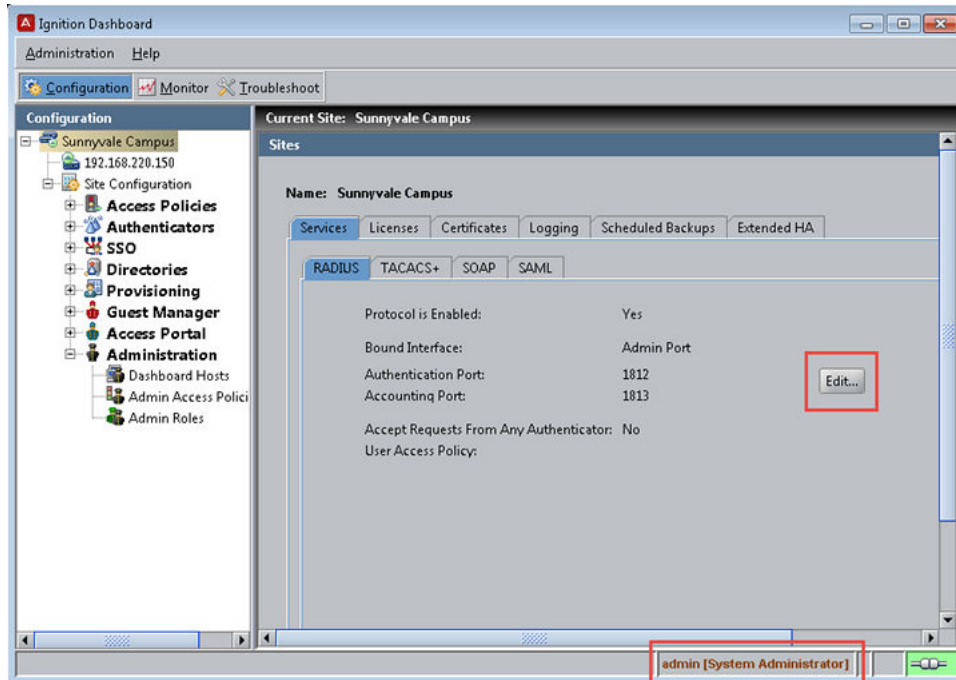
2. Double-click any of the Admin Role Summary entries, or highlight any entry and click **View** to display the groups to which that administrator role belongs. The selected administrator role has all the permissions that are assigned to the groups to which it belongs (Read/Write access).

If the user's administrator role does not belong to a group, only RO (Read Only) access is available.

For example, the System Administrator role has permission to "Site Management", which the other roles do not have. A user who logs in with a non-System Administrator role cannot perform any write operations on the Site Node in the Dashboard hierarchy; only RO access is granted.



For example, a System Administrator has permission to manage the Site Management Functional Group. This means that the System Administrator can perform any create/edit/delete operations under the Site Node. In the following example, the System Administrator can edit the RADIUS configurations in the Services section of Site 0.



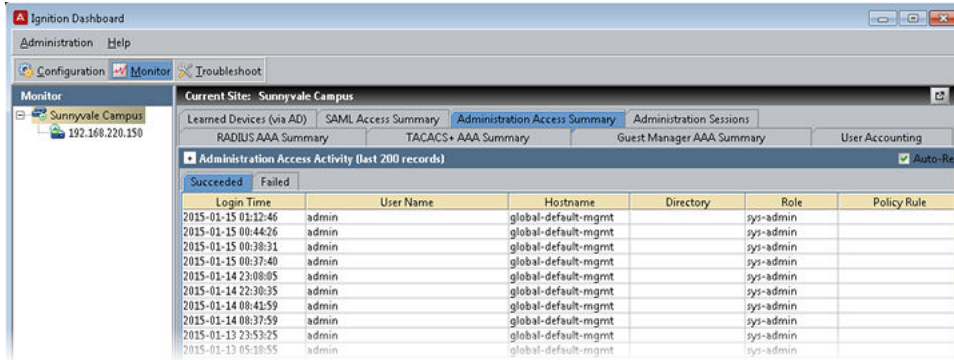
The Monitor Administrator does not have permission to manage the Site Management Functional Group. Only READ functionality is available; the Edit option is disabled (grayed out).

Navigating the Admin logs

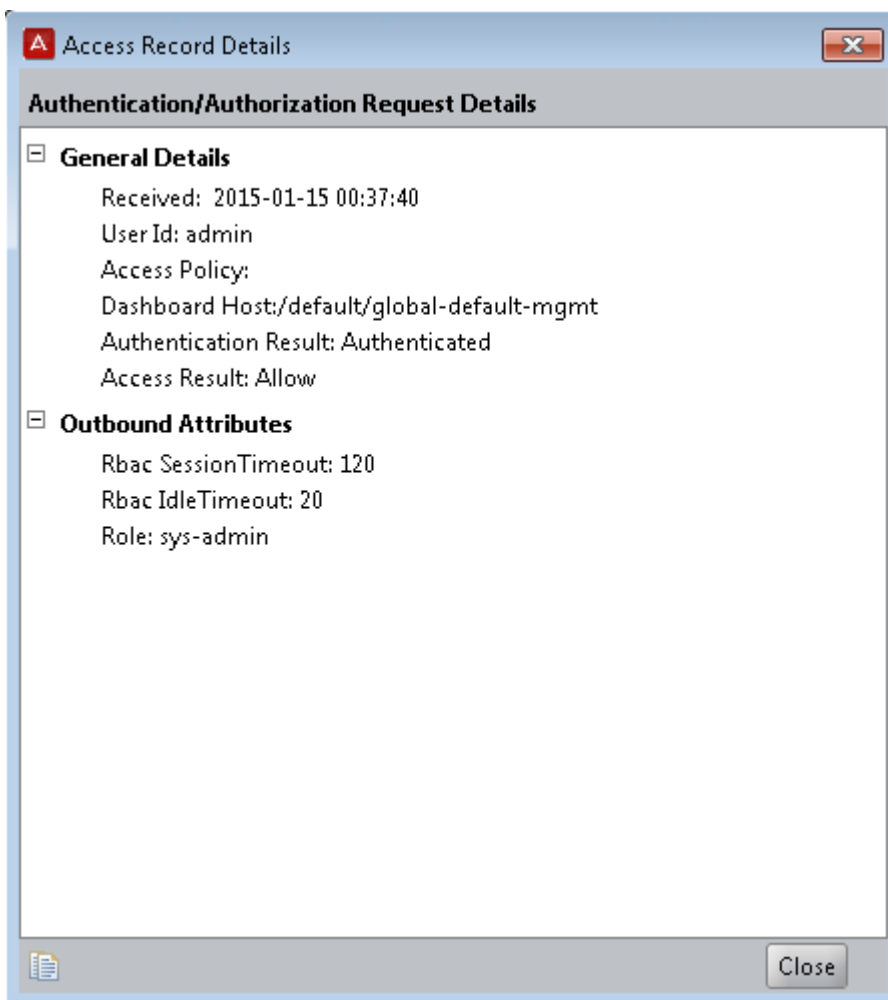
This procedure shows the functionality and information that is available on the Admin Access Policies windows.

Procedure


1. At the top of the main Dashboard window, click **Monitor**.
2. In the **Monitor** hierarchy tree, click your site name.
3. Click the **Administration Access Summary** tab to display a log of the most recent activity performed by the users of your site. For every user who signs in, these logs track information such as login time, user name, host name, directory, role, and policy rule.



- Highlight any entry, right-click it, and select **Record Details** to display more information about the entry.



- Click the **Administration Sessions** tab to display a log of the most recent user sessions on your site. Every user signing in to Ignition Dashboard is associated with a session that tracks information such as the user name, host IP address, role, session start time, and session end time.



Configuring administration preferences

The **Preferences** window allows you to specify Dashboard's time-out settings and log display settings. To open the window, select the command **Administration > Preferences** from the Dashboard main window. See

- [Configuring the idle time-out for Dashboard](#) on page 51
- [Setting viewing preferences for the Monitor view](#) on page 52

Configuring the idle time-out for Dashboard

The Dashboard window locks automatically after a period of inactivity. The default value is 20 minutes. You can turn off the idle time-out function or configure another idle time-out value in the **Preferences** window.

* Note:

If the idle time-out value has been assigned through an Admin Access Policy for a Configuration, Monitor, or Troubleshoot Administrator, this value can not be overridden.

Configure your locking and idle time-out preferences in the **Preferences** window as described in the following procedure.

Procedure

1. Select **Administration > Preferences** in the Dashboard main window.
2. Do one of the following:
 - a. To turn off locking, select **Do Not Lock Ignition Dashboard**.
 - b. To turn on window locking, deselect **Do Not Lock Ignition Dashboard**, and specify the idle time-out period in minutes in the **Wait** field (not available if the idle time-out value has been assigned through an Admin Access Policy). If this option is selected, then after the specified period of inactivity, Dashboard locks and requires a password for unlocking.

To change the password, see [Configuring the System Administrator password](#) on page 57.

Setting viewing preferences for the Monitor view

The **Logging** and **Monitor** tabs of the Preferences window enables you to configure the viewing preferences for Dashboard's Log Viewer tab. For more information on the Log Viewer, see [Viewing and managing logs](#) on page 451.

Follow this procedure to configure your log viewing preferences.

Procedure

1. Select **Administration > Preferences** from the Dashboard main window.
2. Click the **Logging** tab.
 - Select **Automatically refresh logs on tab selection** 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 unselected, then you must use the **Refresh** button in the Log Viewer to load log messages.
 - In the **Order to display log records** section, select **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 select **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. Select **Fit in screen** if you want the Log Viewer to load enough records to fill the window. To set a custom page size, select **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 such as 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 456.

3. Click the **Monitor** tab. The number of authentication/authorization records to display field limits the number of records shown in the site level **AAA 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** shows 200 records, and the **TACACS+ AAA Summary** and **Guest Manager AAA Summary** tabs 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 determine the version of Ignition Dashboard you are running, select **Help > About** from the main window. To determine the firmware version, see [Checking the Firmware version](#) on page 429.

Chapter 5: 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.

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 is important to understand these two concepts: an *Ignition Server site* and an *Ignition Server node*. The *site* is your entire Ignition Server installation; a *node* is an individual Ignition Server appliance.

Depending on your configuration, your Ignition Server *site* can 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:

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

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

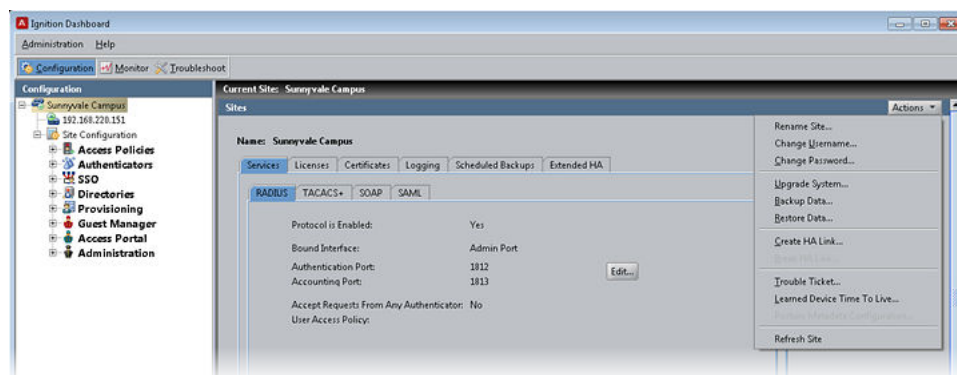
- The **Configuration Hierarchy** navigation panel on the left. Here you specify and organize the nodes in your site.
 - When you select a **site** in the navigation panel, Ignition Server displays the statistics and commands you need to manage the site. See [Managing a site](#) on page 55.
 - When you 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 62.
- A drop-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

The Dashboard **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 37.

For information on managing paired server High Availability sites, see [Paired server high availability configuration](#) on page 391.



Procedure

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 the following.
 - Select a command from the **Actions** menu.
 - 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 56.
- **Change Username:** See [Changing the System Administrator login name](#) on page 56.
- **Change Password:** See [Configuring the System Administrator password](#) on page 57.
- **Update Firmware:** See [Loading a Firmware Image or Package](#) on page 431.
- **Backup Data:** See [Creating a backup](#) on page 422.
- **Restore Data:** See [Restoring from a backup file](#) on page 425.
- **Create HA Link:** See [Run the HA Wizard](#) on page 393.

- **Break HA Link:** See [Breaking an HA pair using Dashboard](#) on page 408.
- **Trouble Ticket:** See [Generating a trouble ticket](#) on page 478.
- **Learned Time to Live (TTL):** See [Setting TTL for Windows Machine authentication](#) on page 325.
- **Refresh Site:** Reloads all site data into Ignition Dashboard.

Renaming an Ignition Server site

Follow this procedure to change the name of an Ignition Server site.

Procedure

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. The default name is Site 0.
2. Choose **Actions Rename Site.....**
The **Rename Site** dialog box displays.
Ignition Server displays the name for the selected site.
3. Enter the new name for the site.
4. Click **OK** to apply your changes.

Changing the System Administrator login name

The default administrator login name is *admin*. Follow this procedure to change the System Administrator login name.

Procedure

1. In Dashboard's **Configuration Hierarchy** tree, click on the name of your site.
2. Right-click and 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**.

Configuring the System Administrator password

The default password is *admin*. Follow this procedure to change the System Administrator password.

Procedure

1. In Dashboard's **Configuration Hierarchy** tree, click on the name of your site.
2. Click **Actions > Change Password**.
3. Enter the existing password in the **Old Password** field.

 **Note:**

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

4. Type the **New Password**. Type the new password again in the **Confirm New Password** field, and click **OK**.

Password Guidelines

Avaya recommends the following guidelines for passwords:

- 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 =.


Managing Ignition Server services

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

- [Configuring Ignition Server's RADIUS service](#) on page 57
- [Configuring Ignition Server's SOAP service](#) on page 59

Configuring 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 405. Use the RADIUS tab to bind the RADIUS service and configure its port numbers.



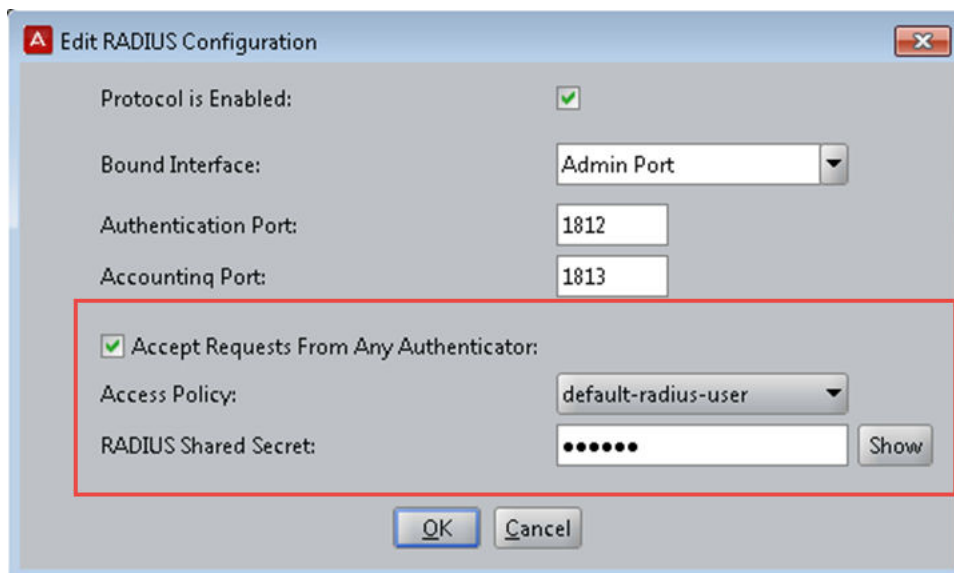
Editing RADIUS communication settings

Follow this procedure to edit RADIUS configuration settings.

Procedure

1. In the Dashboard main window, in the Configuration Hierarchy panel, click the name of your site.
2. In the Sites panel, click the **Services** tab and click the **RADIUS** tab.
3. Click **Edit** in the **RADIUS** tab.

The Edit RADIUS Configuration page displays.



4. Edit as necessary:

- **Protocol is Enabled:** Make sure this check box is selected.
- **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 405.

- **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 454.

5. 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. A global authenticator requires an Ignition Server Base LARGE license. See [Assigning the SOAP service certificate](#) on page 90.

Configuring 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.

This section explains how to configure SOAP API; however, if you are configuring your Guest Manager connection, Avaya recommends that you instead follow the instructions in "Set up Connection to the Ignition Server" in *Configuring Avaya Identity Engines Ignition Guest Manager*, NN47280–501 which describes the additional tasks you must perform in Guest Manager.

Follow this procedure to configure the SOAP service.

Procedure

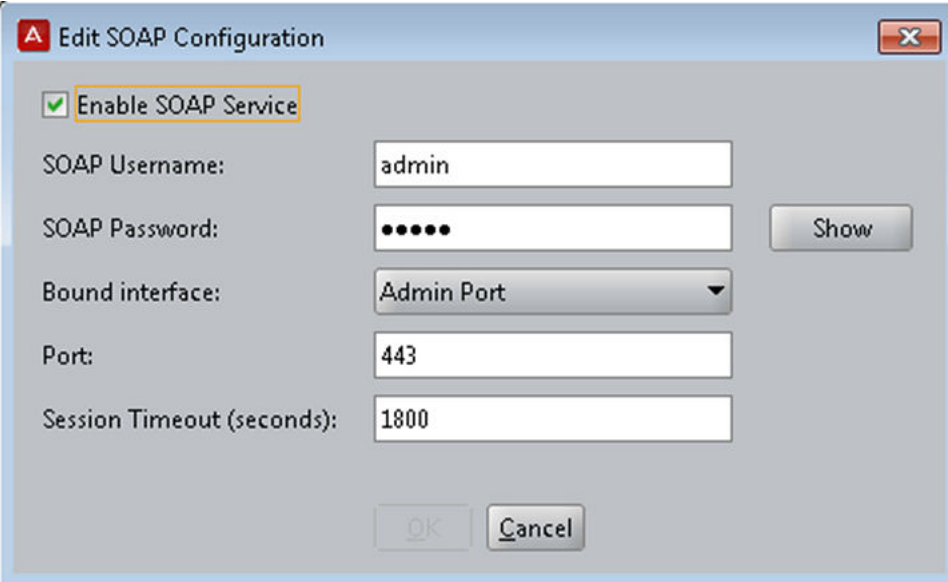
1. In Dashboard's Configuration Hierarchy panel, click the name of your site.

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 [Managing Virtual Interfaces \(VIPs\)](#) on page 405.

3. Click **Edit** in the SOAP tab.

The Edit SOAP Configuration window displays.



4. Configure the SOAP connection parameters:
 - a. **Enable SOAP Service** — Select this check box to make the SOAP API service available.
 - b. **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.
 - c. **SOAP Password** — Password that the SOAP user account uses to connect.
 - d. **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 405.
 - e. **Port** — Enter the port number to which API clients should connect. Traffic through this port is HTTPS traffic.
 - f. **Session Timeout** — Enter the time 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 time-out period, but it does not cause Guest Manager to become disconnected since Guest Manager automatically reconnects.

! **Important:**

Configure the SOAP **Session Timeout** to a period of 180 seconds or longer. Configuring it as 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 described in [Managing Virtual Interfaces \(VIPs\)](#) on page 405.
7. Perform SOAP configuration steps in Guest Manager.
 - a. Install a copy of the SOAP certificate in Guest Manager as explained in *Configuring Avaya Identity Engines Ignition Guest Manager*, NN47280–501 in the section, “Installing a SOAP Certificate.”
 - b. Make SOAP and RADIUS settings in Guest Manager as explained in *Configuring Avaya Identity Engines Ignition Guest Manager*, NN47280–501 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. Use the following procedure to reset the SOAP password.

Procedure

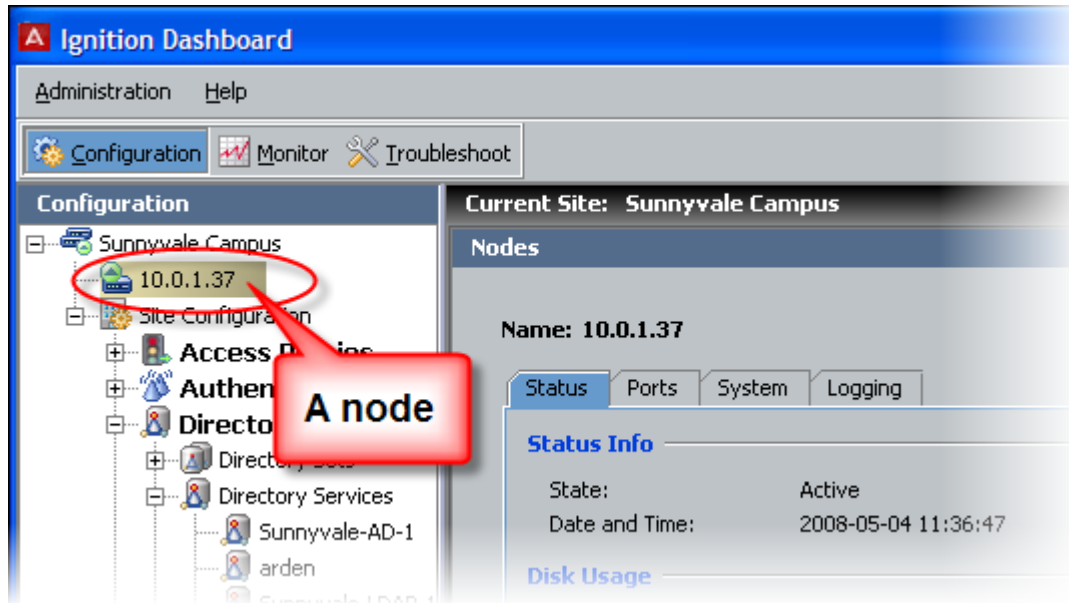
1. In Dashboard’s Configuration Hierarchy panel, click the name of your site.
2. In the Sites panel, click the **Services** tab and click the **SOAPS** tab.
3. Click **Edit** in the **SOAP** tab.

The Edit SOAP Configuration window displays.
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.



Actions menu for a node

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

- Reboot
- Power Down
- Reinitialize
- View Logs
- Rename Node

Rebooting a node

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

Procedure

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**. Alternatively, select **Actions > Reboot**.

The Reboot Confirmation window displays, requiring you to confirm your action.

3. Click **Yes**.

Ignition Server disconnects the node and reboots the Ignition Server.

4. 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.

Procedure

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**. Alternatively, select **Actions > Power Down**.

The Power Down Confirmation window displays, requiring you to confirm your action.

3. 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 Ignition Server from Dashboard

Ignition Server allows you to reinitialize a node only when the selected node is a standalone node.

Important:

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.

Procedure

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:
 - a. In Dashboard's Configuration tree, click the name of your site.
 - b. Click the **Licenses** tab and click **Export All KRS Licenses**. 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 tree, click the name or IP address of your node.
4. Right-click on the selected node and choose **Reinitialize**. Alternatively, select **Actions > Reinitialize**.

A confirmation window appears, requiring you to confirm your action.
5. Click **Yes** to proceed with the reinitialization.

Ignition Server resets the selected node to its factory settings.
6. After the Ignition Server has rebooted, configure its IP address.
7. 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 78.

Viewing logs for a node

Follow this procedure to view the logs of your Ignition Server node.

Procedure

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**. Alternatively, select **Actions > View Logs**.

This opens the Monitor tab of Dashboard and places you in the Log Viewer tab for your node.

Renaming a node

Follow this procedure to rename your Ignition Server node.

Procedure

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**. Alternatively, select **Actions > Rename Node**.
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 a selected node. 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: 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

- **Software Version:** indicates the version and build number for the firmware on the Ignition Server.
- **Model:** indicates the VM model.
- **Installation Date:** indicates the installation date.
- **Last Boot Date:** indicates the last time the node was rebooted.
- **Image Creation Date:** indicates the date that the system image was created.
- **Serial Number:** displays the 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 75.

Hypervisor Information: provides information about the system hypervisor type, vendor and version.

Obtaining the Ignition Server Serial Number

The Avaya Identity Engines Ignition Server software ships without any licenses. There are seven different software licenses that can be installed on Ignition Server: Base License, Guest Manager License, NAP Posture License, TACACS+ License, Ignition Reports License, Access Portal License, and Avaya Aura® Single-Sign-On (SSO) License. At a minimum, you must obtain the Base License to be able to configure and run the server.

If you are applying a NAP Posture License or an Access Portal License, select the Access Portal License that matches the Ignition Server Base License (LITE, SMALL, or LARGE).

* Note:

Beginning with Identity Engines Release 9.0, Identity Engines starts to transition from DVD delivery to electronic software delivery. Depending on how you place your order, you may

receive DVDs with paper LACs, or electronic software delivery and electronic LACs. With each method you will receive instructions on how to obtain your licenses.

Once you have purchased Identity Engines, depending on how you placed your order you receive either a set of DVDs accompanied with paper LACs (License Authorization Codes), or else you receive electronic delivery of your LAC by email and you then download the software from Avaya support site via PLDS.

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/>.

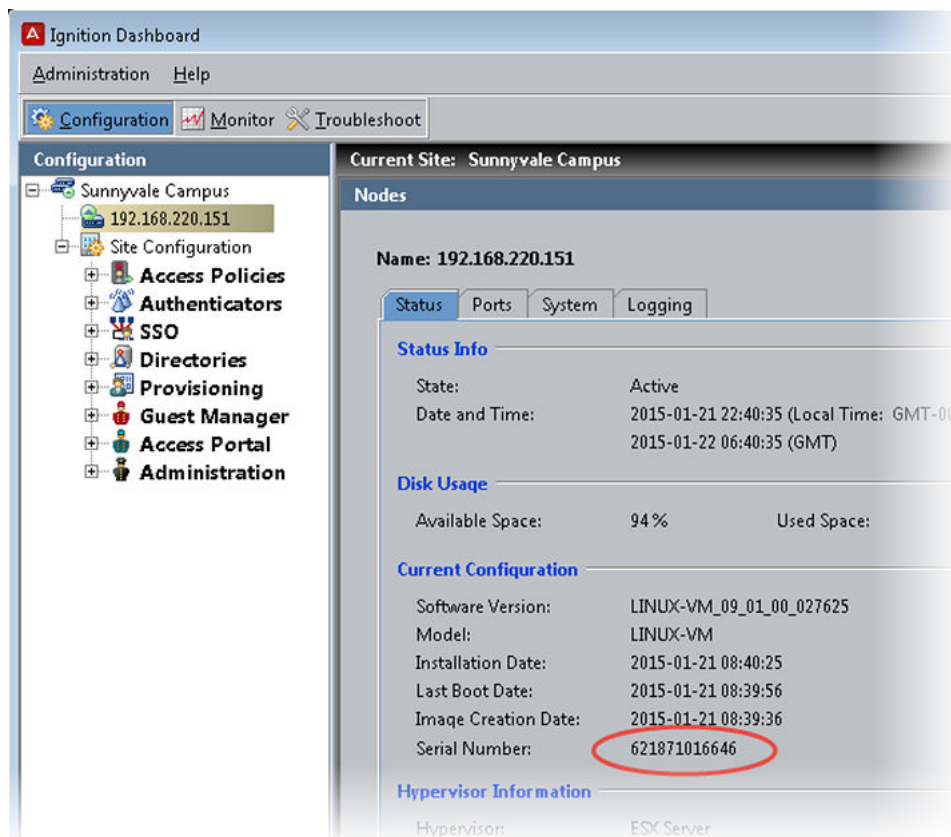
Once you have installed both the Ignition Server Virtual Appliance and the Ignition Dashboard, you must obtain the Ignition Server node Serial Number (also known as the Host-ID) from the Dashboard. The Ignition Server Serial Number is required in order to generate licenses regardless of whether they are KRS licenses or PLDS licenses. Beginning with Release 9.0, the Ignition Server Serial Number is always a string of 12 digits.

If you have a paired server High Availability (HA) deployment, you need to obtain the Serial Numbers of both Ignition Servers that make up the HA-pair.

Procedure

Do one of the following:

- In the VMWare vSphere Client, launch the Ignition Server CLI and enter the command `show version`.
- From the Dashboard Configuration tree, click the name or IP address of your node, click the **Status** tab and note the serial number.



System tab

The **System** tab displays information about the system operating system, enables you to specify the DNS servers, configure the routing information, and configure time synch, SNMP, SSH, and SMTP settings for the node.

Viewing Ignition Server's DNS settings

Procedure

1. In Dashboard's Configuration tree, click the name or IP address of your node.
2. Click the **System** tab , and then the **DNS** tab.

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.

Procedure

1. In Dashboard's Configuration tree, click the name or IP address of your node.
2. Click the **System** tab and then the **DNS** tab.
3. Click **Edit**.
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 and system routing tables. To add a network route, see [Adding a route to Ignition Server's Routing Table](#) on page 69. To edit a network route, [Editing an existing route](#) on page 69. To delete a network route, see [Deleting an existing Route](#) on page 70.

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 you configure routes in the **Static Routing** configuration table, make sure you have configured the IP addresses of the Ignition Server interfaces you plan to use.

Adding a route to Ignition Server's Routing Table

Follow this procedure to add a network route to Ignition Server's Static Routing table.

Procedure

1. In Dashboard's Configuration tree, click the name or IP address of your node.
2. Click the **System** tab and click the **Static Routing** tab.
3. Click **Add**.

The Add a Route window displays.

4. Add a gateway:
 - **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**.

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

Editing an existing route

Follow this procedure to edit a network route in Ignition Server's Static Routing table.

Procedure

1. In Dashboard's Configuration 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. Right-click the entry and click **Edit**.
The Edit a Route window displays.
5. Edit the route as required.
6. Click **OK**.
The Static Routing configuration list shows the updated entry for the route.

Deleting an existing Route

Follow this procedure to delete an existing network route.

Procedure

1. In Dashboard's Configuration tree, click the name or IP address of your node.
2. Click the **System** tab and click the **Static Routing** tab.
3. 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 [Problem: Authentication fails on Active Directory](#) on page 484.

SSH settings

You can configure an SSH network port on the Ignition Server and connect to the Ignition Server CLI through SSH. See [Managing Ignition Server licenses](#) on page 75.

SMTP settings

You can configure Ignition Server to send log alerts through e-mail using an SMTP server on your network. See [Sending log messages Via E-Mail](#) on page 444.

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 following table explains what sort of traffic each interface can carry.

	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
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:

- **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 to which the port is connected. In most of these connections, this link is up regardless of the physical port to which it is tied. Only if the VM-HOST disconnects the port, does the LinkStatus 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 configure.

Field Name	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.
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 in this table) and check your network connections and cables.
IP Address	Any valid IP address	The IP Address field in the Edit PortConfiguration 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 IP address is configured 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.

Follow this procedure to change the IP address settings for the Admin Port.

Procedure

1. In Dashboard's Configuration tree, click the name or IP address of your node.

2. Click the **Ports** tab and click the **Admin Port** entry.
3. Click **Edit**.

The Edit Port Configuration window displays the current Admin Port IP address.

4. Change the IP address by entering the new IP address and mask.
5. Click **OK**.

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.

6. 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. To reconnect Dashboard to the Ignition Server, 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 **Edit**, Dashboard displays the **Edit Port Configuration** window. Here you can enable or disable a service port and configure the IP address for the port.

Enabling the service port

Follow this procedure to enable the service port and set its IP address.

Procedure

1. In Dashboard’s Configuration tree, click the name or IP address of your node.
2. Click the **Ports** tab and click the **Service Port** entry.
3. Click **Edit**.

The Edit Port Configuration window displays.

By default, the **Enable Port** check box is not selected, indicating that the port is disabled.

4. To enable the port, select the **Enable Port** check box.
5. In the **IP Address** fields, enter the IP address and subnet mask entries for the port.

No two port IP addresses can be located on the same subnet.

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

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

7. Assign services to the port or ports you enabled. See [Managing Ignition Server services](#) on page 57.

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 configure the IP address and other settings of your HA port. You can also configure the HA port IP address by following the instructions detailed in the section [Enabling the service port](#) on page 73.

Enabling the HA port

Follow this procedure to enable the HA port and set its IP address.

Important:

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

Procedure

1. In Dashboard's Configuration tree, click the name or IP address of your node.
2. Click the **Ports** tab, and click the **HA Port** tab.
3. Click **Edit**.

The Edit Port Configuration window displays. By default, the **Enable Port** check box is not selected, indicating the port is disabled.

4. To enable the port, select the **Enable Port** check box.
5. In the **IP Address** fields, enter the IP address and subnet mask entries for the port.
6. Click **OK** to apply your changes.

Managing Ignition Server licenses

The Avaya Identity Engines Ignition Server software ships without any licenses. There are seven different software licenses that can be installed on Ignition Server: Base License, Guest Manager License, NAP Posture License, TACACS+ License, Ignition Reports License, Access Portal License, and Avaya Aura® Single-Sign-On (SSO) License. At a minimum, you must obtain the Base License to be able to configure and run the server.

If you are applying a NAP Posture License or an Access Portal License, select the Access Portal License that matches the Ignition Server Base License (LITE, SMALL, or LARGE).

*** Note:**

Beginning with Release 9.0, Identity Engines starts to transition from DVD delivery to electronic software delivery. Depending on how you place your order, you may receive DVDs with paper LACs, or electronic software delivery and electronic LACs. With each method you will receive instructions on how to obtain your licenses.

This section explains how to install and manage licenses.

Checking Ignition Server licences

Whenever you log in to the Dashboard, a license validity check activates to check for validation and expiry date of the license.

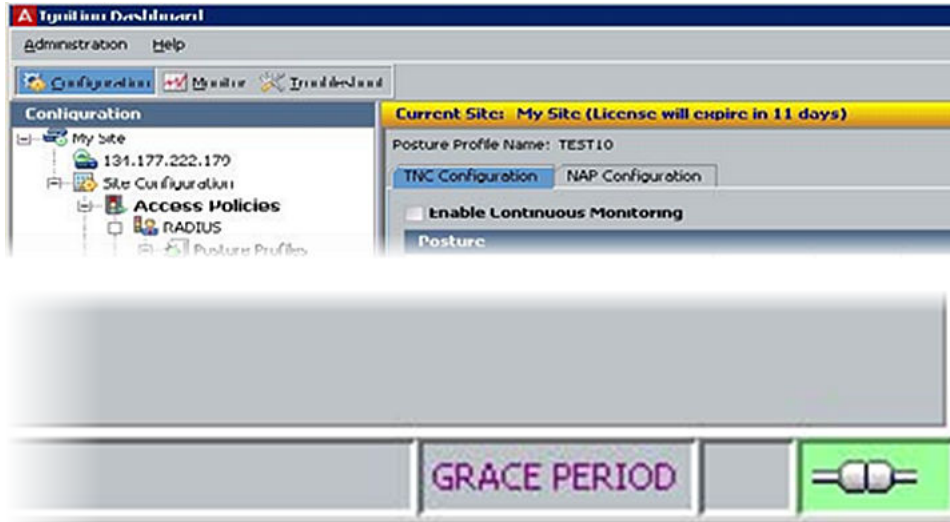
If the license is not valid or is expired, a message displays to prompt you to enter a valid license. This check includes the validity of a 30-day grace period that may be provided during upgrading. As shown in the following figure, if the grace period expires then the Ignition Server stops processing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

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 the 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.



Temporary 30-day licenses can be obtained from <http://www.avaya.com/identitytrial>.

Follow this procedure to check your Ignition Server licenses.

Procedure

1. In Dashboard Configuration tree, click the name of your site.
2. Click the **Licenses** tab.

The Licenses list shows the installed licenses.

3. Click on a license to see the license details.

The **Valid From** and **Valid Until** fields 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. To email or copy the license details, click **Copy to Clipboard**, and paste it into your email or other application.

Seat limit enforcements

As of Identity Engines Release 9.0.3, the three different levels of Ignition Server, based on the number of authenticators allowed, are supported as follows.

- Ignition Base LARGE - Unlimited authenticators
- Ignition Base SMALL - 20 standard authenticators + 300 x WLAN 9100 APs
- Ignition Base LITE - 5 authenticators + 75 x WLAN 9100 APs

Seat limit enforcement occurs 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 are automatically set to a 'disabled' state. You 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, then 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. You can selectively enable the authenticators.

Obtaining PLDS licenses

If you have received your LAC by electronic delivery (email), your licenses are PLDS licenses.

Using the Avaya Product Licensing and Delivery System (PLDS), you can activate the license entitlements and download the products.

Upon your purchase of Identity Engines, you receive an electronic LAC with which you, as a customer or Avaya Business Partner who has permissions in PLDS for your site or sales order, can access PLDS and generate license entitlements for you. You must provide the Serial Number, also known as the host ID, of the Identity Engines Ignition Server and your LAC in order to generate a license. The LAC helps you to identify the product among other Avaya products you hold licenses for and to keep track of the number downloads, while keeping the required groups and coordinators informed through e-mail messages. The LAC e-mail recipients must be identified during the order placement process by providing their e-mail addresses.

With the LACs in hand, you can use the Quick Activation screen to activate the LACs and download the product.

Obtaining KRS licenses

If you received paper LACs with your purchase, follow the instructions on the paper LACs regarding how to obtain your licenses. These will be KRS licenses.

Send an email to datalicensing@avaya.com to request your KRS licenses and include the following information:

1. End user company name and full mailing address (no mailboxes).

2. End user company URL.
3. End user contact name.
4. End user corporate email address.
5. End user phone number.
6. License Authorization Code (LAC) that shows in the box at the bottom right of the LAC certificate.
7. Serial Number or Serial Numbers if you have an HA deployment.

After the information is verified, licenses are sent to you by email.

Installing an Ignition Server license

Avaya Identity Engines currently supports the KeyCode Retrieval System (KRS) based and the Avaya Product Licensing and Delivery System (PLDS) licensing models.

The Avaya PLDS provides customers, Business Partners, distributors and Avaya Associates with easy-to-use tools for managing asset entitlements and electronic delivery of software related licenses. Using PLDS, you can perform activities such as license activation, license de-activation, license re-host, and software downloads.

There are a few key differences between the two types of licenses which are important to understand, especially if you will be using both types of licenses.

Important:

Note the following:

- Beginning with Release 9.0, Ignition Server supports both KRS and PLDS licenses to accommodate customers who do not yet have access to Avaya PLDS. Over time, Identity Engines will transition to support a single licensing system — PLDS.
- An important difference between KRS licenses and PLDS licenses is that KRS licenses are individual licenses, while a PLDS license file always includes all PLDS licenses within a single PLDS license file, which is in XML format.
- A PLDS license file ALWAYS has, at a minimum, a Base license.
- KRS licenses can be exported from the Dashboard and saved on your desktop.
- PLDS licenses cannot be exported from the Dashboard. Therefore, it is very important to ALWAYS safeguard the PLDS license file you have received from PLDS. You may be able to log back in to PLDS and regenerate the license file again.
- Installing PLDS licenses deletes any PLDS and KRS licenses that are already installed . Therefore, it is important to export all KRS licenses before installing PLDS licenses in order to safeguard your KRS licenses.
- Since KRS licenses are deleted when installing PLDS license, then before installing PLDS license you MUST export and save the KRS licenses if any already exist.

- Installing KRS licenses overwrites any installed PLDS licenses.

Procedure

1. In the Dashboard's Configuration tree, click the name of your site and click the **Licenses** tab.
2. Click **Install**.
 - a. Browse to the license file location, select the appropriate file, and click **OK**.
OR
 - b. 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. If it is a KRS license, copy the whole license including "BEGIN IGNITION LICENSE CERTIFICATE" and "END IGNITION LICENSE CERTIFICATE". If it is a PLDS license, copy the entire XML file.

Return to the License Installation window of Dashboard and click **Paste** to paste the license text there.
 - c. Click **OK**.

Replacing an Ignition Server license

Warning:

Before you delete your old license, make sure you have obtained its replacement.

Follow this procedure to replace a license.

Procedure

1. In the Dashboard's Configuration tree, click the name of your site and click the **Licenses** tab.
2. Click on the license and click **Delete**.
3. Install the replacement as shown in [Installing an Ignition Server license](#) on page 78.

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.

Procedure

1. In the Dashboard's Configuration tree, click the name of your site and click the **Licenses** tab.
2. Click **Export All KRS Licenses** 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.

Troubleshoot tab

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

- [Running a ping test](#) on page 80.
- [Running a packet capture](#) on page 81.

Procedure

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.

Follow this procedure 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.
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 **Browse** 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 **Stop**.

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

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

Chapter 6: 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.

Required types of Certificates

- The **admin certificate** secures Ignition Dashboard-to-Ignition Server communications. See [Admin certificate](#) on page 86.
- The **SOAP service certificate** secures Guest Manager-to-Ignition Server communications. See [Assigning the SOAP service certificate](#) on page 90.
- **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 91.
- **Protocol root certificates** are used to verify supplicants’ certificates during EAP-TLS and PEAP/EAP-TLS authentication. See [Installing protocol root certificates](#) on page 92.

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 86) or to secure authentication transactions (see [Assigning protocol credential certificates](#) on page 91).
- The **Certificate Requests** tab is used to generate certificate requests. See [Getting a new certificate](#) on page 84.
- 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 92.
- 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 92.

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

- The **Root Certificates** window (open it by selecting **Administration > Root Certificates**) 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 87.

Getting a new certificate

Create the certificate request

Follow this procedure to create the new certificate request.

Procedure

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.



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 key.

! Important:

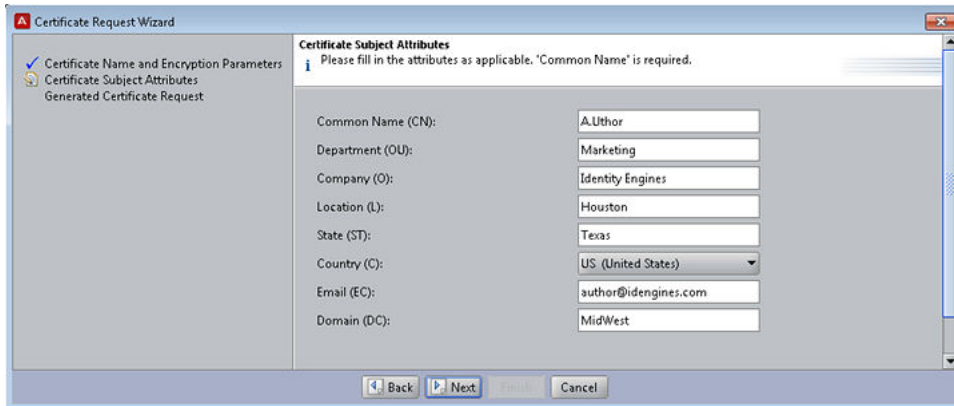
If this certificate is to be used as your Ignition Server **admin certificate**, it must use an RSA key; you cannot use a certificate based on a DSA key for an Ignition Server **admin certificate**.

! Important:

If this certificate is to be used as a tunnel certificate that supports Windows XP clients, observe the limitations explained in [Factors that limit your choice of a Protocol Credential Certificate](#) on page 244.

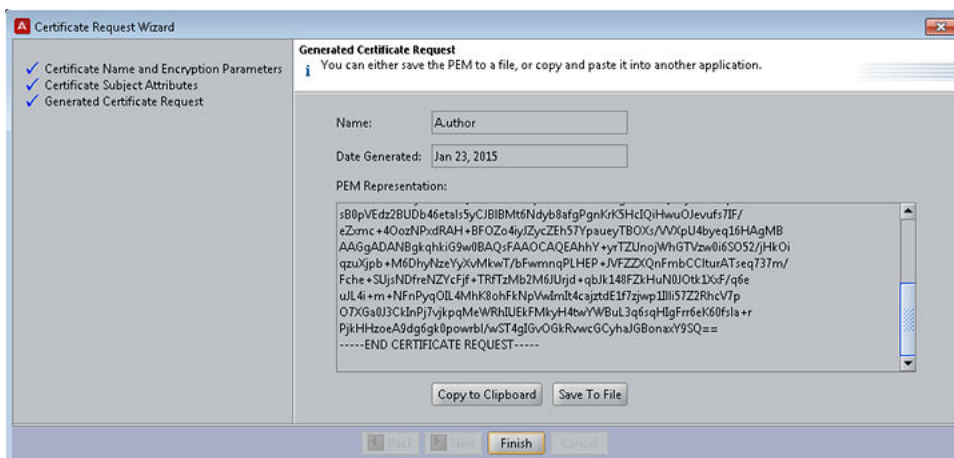
4. Click **Next**.

The Wizard displays the Certificate Subject Attributes window. The **Common Name** is required.



5. Click **Next**.

Ignition Server generates the certificate request and displays the Generated Certificate Request window.



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 [Import the certificate](#) on page 85.

Import the certificate

After you have received the certificate you requested in Step 7 of the previous procedure, or if you have a certificate ready for import, import the certificate as described in the following procedure.

Procedure

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 **Import Certificates**.
4. Navigate to find your certificate.

Make sure that your certificate meets the certificate file requirements listed in [Format of certificate files](#) on page 83.

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 86.

Important:

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.

- Dashboard-Server communications: see [Replacing the Admin certificate](#) on page 88.
- Guest Manager-Server communications: see [Assigning the SOAP service certificate](#) on page 90.
- User authentications: see [Assigning protocol credential certificates](#) on page 91.

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**. Replace the default certificate as soon as possible after installing Ignition Server. See [Replacing the Admin certificate](#) on page 88 for instructions.

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 following procedure explains 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, see [Replacing the Admin certificate](#) on page 88.

! Important:

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.


Follow this procedure to install a copy of the admin certificate on Dashboard.

Procedure

1. Contact your System 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 83.

2. From the Dashboard main window, select **Administration > Root Certificates**.
3. In the **Root Certificates** window, click **Add**.



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 following procedure explains 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 in Step 1.

Follow this procedure to replace the admin certificate.

Procedure

1. If you have not yet requested or imported your admin certificate into Ignition Server, do so as explained in [Adding a certificate revocation list URL](#) on page 92.
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!*)
 - 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:

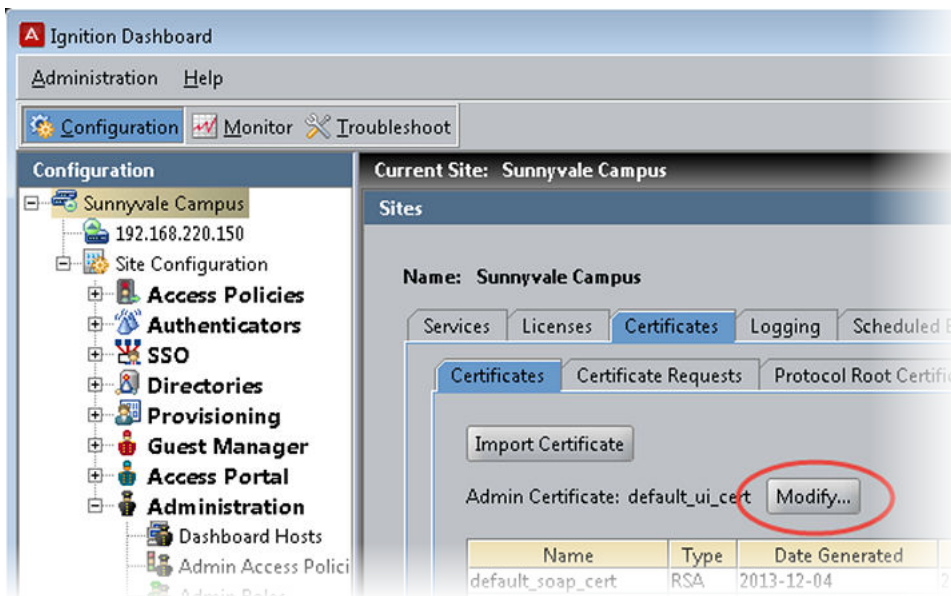
If you choose an alias that is already in use, the newly imported certificate replaces the certificate previously aliased under that name. *Do not replace a certificate that is still needed for communicating with one of your Ignition Servers!* If you do so, you cannot connect to that Ignition Server.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

- Click **Add**.

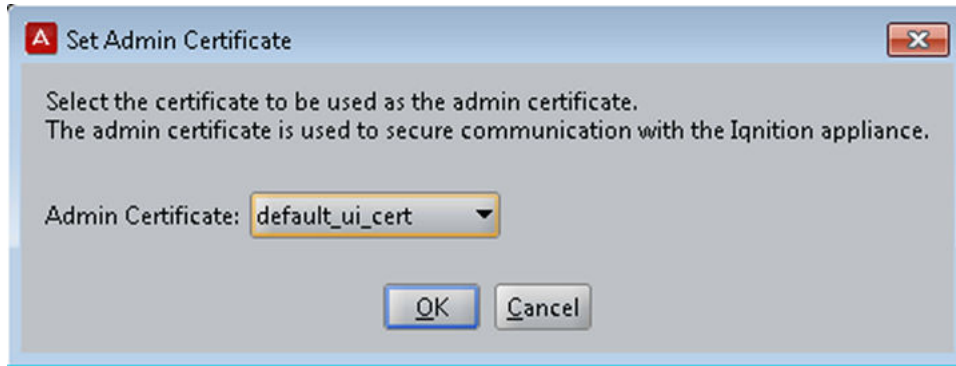
Ignition Server adds the selected entry to the display in the **Root Certificates** list.

- 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**.



7. Click **OK**.

Ignition Server displays a confirmation window. If you performed Step 2 in this procedure, 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 **Administration > Root Certificates** from the Dashboard main window, select the certificate in the list, and click **Delete**. *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 following procedure explains 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.

Procedure

1. If you have not yet imported your certificate into Ignition Server, do so as explained in [Adding a certificate revocation list URL](#) on page 92.
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 **Modify**.
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 Guest Manager Configuration*, NN47280-501.

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 the corresponding root certificate installed on it, 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 that 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

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 [Adding a certificate revocation list URL](#) on page 92.
3. In the Dashboard’s Configuration tree, expand **Access Policies > RADIUS**, and click on the name of your access policy.
4. Click the **Authentication Policy** tab and click **Edit**.
5. In the **Authentication Policy** window, go to the **Protocol Credential** section.
6. In the **Certificate** drop-down list, select the name of your protocol credential certificate.
7. Click **OK**.

(See [Creating an authentication policy](#) on page 244 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.

Procedure

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 83.
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, go to [Adding a certificate revocation list URL](#) on page 92 to configure your Certificate Revocation Lists.

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 94.)

Adding a Certificate revocation List URL to Ignition Server

Follow this procedure to add a certificate revocation list URL.

Procedure

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 **New** 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 must request a new certificate and, if necessary, update the list of URLs stored in Ignition Server.

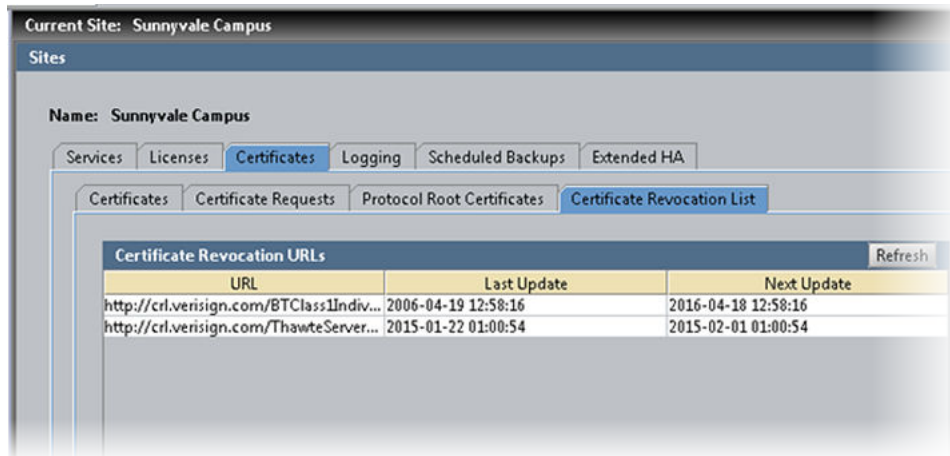
Viewing the certificate revocation list URLs

Follow this procedure to view the URLs of the certificate revocation list.

Procedure

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 of an entry under this tab.



URL

Each entry displayed in the **URL** 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 **Refresh**, 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.

CLI Command to toggle CRL check level

In a multi-layer Certificate Authority (CA) configuration, a certificate revocation list (CRL) check can happen at the certificate-issuing CA node or up at the root CA node. CLI commands are provided to configure the CRL check level.

Command format	Definition
<code>radius crl leaf</code>	Check CRL at the certificate-issuing CA node.
<code>radius crl root</code>	Check CRL at the root CA node.

Viewing a certificate

Follow this procedure to view a certificate.

Procedure

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 **View**.
Ignition Server displays the contents of the selected certificate.
5. Click **OK**.
6. To view the copy of the admin certificate saved in Dashboard, select **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.)

Follow this procedure to delete a certificate.

Procedure

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.

Deleting an existing certificate request renders unusable any certificate the CA sends you based on this request.

Ignition Server prompts you to confirm the deletion request.

5. 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

Follow this procedure to view an existing certificate request.

Procedure

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.

Chapter 7: Authenticators

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

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 configure 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.

1. *Authenticator bundles* allow you to represent all authenticators on a subnet with a single authenticator record
2. 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 104 for information on bundles, and see [RADIUS global authenticator](#) on page 112 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 in this section.

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*.
- Next it tries to match *small authenticator bundles* (large netmask).
- Next it tries to match *large authenticator bundles* (small netmask).
- Finally, it tries to match 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. Therefore, 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 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.

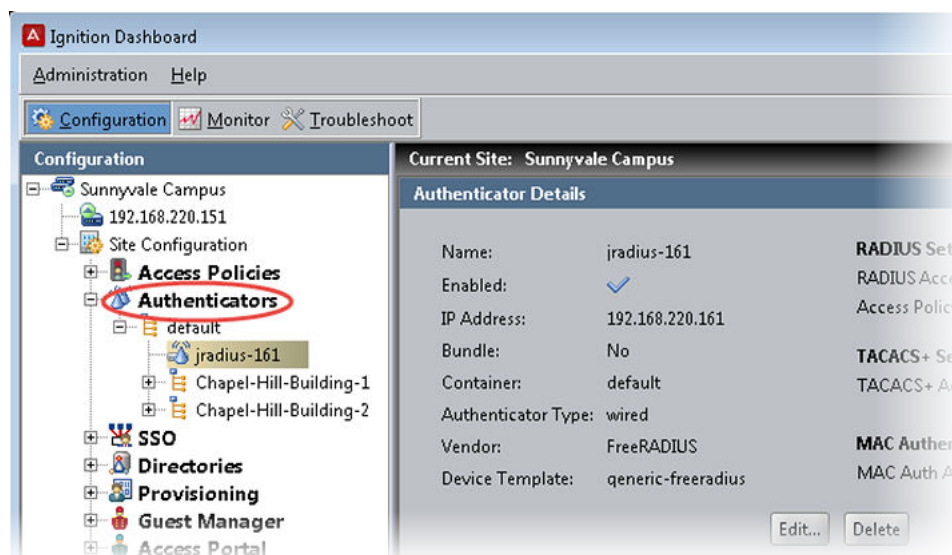
A container holds a set of authenticators you have grouped in Ignition Server, and the authenticator hierarchy is the tree of these containers. Organizing containers into a tree allows you to group your authenticators (for example, geographically) for ease of management, and allows you to create authorization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

Each authenticator belongs to exactly one container and has exactly one access policy. A container can 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 101. 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 110.



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 configured 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 246.

- **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 260.

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 using any authenticator that is in the *Colorado* container. (Expressed in Ignition Server's rule-writing terminology, 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*.)

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.

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 the **Include descendants of selected container** check box is selected, it also lists the authenticators associated with all sub-containers of the selected container. When this check box 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

Configure 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.

Follow this procedure to create the hierarchy.

Procedure

1. In the Dashboard's Configuration tree, expand **Site Configuration > Authenticators**, and click on the container that you want to be the root of your hierarchy. Typically this is the "default" container.
2. Select **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 selecting **Actions > Add Container** 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 **New** 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 110.

Adding a container to an authenticator hierarchy

Follow this procedure to add a new container to the authenticator hierarchy.

Procedure

1. In the Dashboard's Configuration tree, expand **Site Configuration > Authenticators**, and click on the parent container under which you are defining a new container.
2. Right-click the parent container and select **Add Container**. Alternatively, select the parent container, and select **Actions > Add Container**.

The Add Container window displays, listing the name of the parent container to which the new container is to be added.

3. Enter the unique name for the new container in this window.
4. 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 482.

Procedure

1. In the Dashboard's Configuration tree, expand **Site Configuration > Authenticators**, and click on the container.
2. Right-click on the container and select **Rename Container**. Alternatively, select **Actions > Rename Container**.
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 tree in the main window.

 **Note:**

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 482.

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.

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

Procedure

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 118.)
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 tree, expand **Access Policies** > **RADIUS**, and 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 **Edit** 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 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 tree, expand **Site Configuration** > **Authenticators**. Find your authenticator in the tree.
5. Right-click your authenticator and select **Delete Container**. Alternatively, select the container, and select **Actions** > **Delete Container**.
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.

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 107 for more information.

If you are using Access Portal as an authenticator, use this procedure: [Registering Access Portal with the Ignition Server](#) on page 108.

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.

Procedure

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(a). Authenticator**. Ignition Server displays the 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 check box to enable the new authenticator.
- **IP Address:** The IP Address of the authenticator.
- **Bundle:** Select this check box 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** check box, 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 to what part of your organization it belongs. (See [Authenticator hierarchy and containers](#) on page 98.) You can change the container association by clicking the blue text. See [Placing an authenticator in the authenticator hierarchy](#) on page 110.
- **Authenticator Type:** Specify what type of device the authenticator is: Any, Wired, Wireless, VPN, SIP, or Other. The 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 291.)

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.

! Important:

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 sub-authenticators](#) on page 114.

 - **Enable MAC Auth:** Select this check box to provide authorization based on the MAC address of the device that is trying to connect. See [MAC Authentication](#) on page 339.
5. If this authenticator will use TACACS+ authentication, click the **TACACS+ Settings** tab and configure the following:
 - Select the **Enable TACACS+ Access** check box. You must select this check box 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 333.
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 configure 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
```

To determine 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 configuration is complete. If you do not already have appropriate security policies defined for the switch, see [User authentication policy](#) on page 237.

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.

Follow this procedure to import several authenticators.

Procedure

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-radiususer, , , , 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, and so on.

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 **Browse**, navigate to your csv file, click it, and click **Open**.

The **Authenticators Import Status** window displays. 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

Follow this procedure to export a set of authenticators.

Procedure

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**.

The **Export Authenticator** window displays the following message:

```
Do you want to export the Authenticators from the child containers?
```

3. Click **Yes** to export the Authenticators from the child containers or click **No** to not export the Authenticators from the child containers

The **Save** window displays.

4. 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 with the Ignition Server

Follow this procedure to register Access Portal with the Ignition Server.

Warning:

Any mismatch in RADIUS configuration between the Ignition Server and Access Portal (for example, server IP address, shared secret, password, and so on) can result in fatal or internal errors to the clients. Always perform a test user authentication after configuring RADIUS settings in Ignition Server and Access Portal.

Procedure

1. In the Dashboard **Configuration** tree, expand the **Access Portal** folder and click **Access Portal Servers**.
2. Click **New**.

The **Access Portal Server Details** page displays.

The screenshot shows the 'Access Portal Server Details' configuration window. It includes the following fields and options:

- Name:** A text input field with a red 'x' icon.
- IP Address:** A text input field with a red 'x' icon.
- Trust Device Update**
- Expiration**
 - Date:** 2014-12-28 09:28:19 (with a clock and calendar icon)
 - Duration:** Days: 30, Hours: 0
- Delete On Expiry**
- RADIUS Shared Secret:** A text input field with a red 'x' icon and a 'Show' button.
- RADIUS Access Policy:** A dropdown menu showing 'default-radius-user'.
- Enable MAC Auth**
- Access Policy:** A dropdown menu showing 'default-radius-device'.
- Authentication methods:
 - Do Not Use Password**
 - Use RADIUS Shared Secret As Password**
 - Use This Password:** A text input field with a 'Show' button.
- Member Of Groups:** A section with a table header 'Internal Group Name' and an empty table body. Below the table are 'Add...' and 'Remove' buttons.

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. Ensure that you enter the IP address of the ADMIN interface. Also make sure that Access Portal's ADMIN interface is reachable from the Ignition Server.
 - **Trust Device Update:** Select this check box if you want the Ignition Server to create a device record in the local store with the device fingerprint of the authenticating user. Note that, 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:** Select this check box if you want to specify an expiry date or lapse period for the devices learned through Access Portal.
 - To specify an expiration date, click **Date** and 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.
 - To specify a lapse period, click **Duration** and use the arrow keys to specify the number of days and hours from the time the device was learned until the time it expires.

- **Delete On Expiry:** Select this check box if you want the Ignition Server to delete learned devices after the expiry date. Note that it may take up to 24 hours after the expiry date for the devices to be purged from the local store.
- **RADIUS Shared Secret:** Enter the Shared Secret that you configured for RADIUS server.
- **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-radiususer).
- **Enable MAC Auth:** Select this check box to provide authentication based on the MAC address of the device that is trying to connect.
- **Member of Groups:** Select one or more internal groups to which unregistered devices can be auto-associated. Click **Add**, select the internal group or groups, and click **OK**.

4. Click **OK**.

The **Access Portal Server Summary** page displays.

Current Site: Sunnyvale Campus					
Access Portal Server Summary					
Server Name	IP Address	RADIUS	RADIUS Access Policy	MAC Auth	MAC Auth Access Policy
CP1	172.15.1.100	✓	Sunnyvale-RADIUS-policy	✗	

Related Links

[Using Ignition Access Portal as a Policy Enforcement Point](#) on page 120

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 configure the authenticator hierarchy (“Container”) label, use the following procedure.


Procedure

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

! Important:

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, select the **Include Selected Hierarchy Descendents** check box.

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.



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

Finding an authenticator

Follow this procedure to find an authenticator record.

Procedure

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 select the **Include Selected Hierarchy Descendents** check box.
 - *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, select the **Include Selected Hierarchy Descendents** check box.

Pinging an authenticator

Follow this procedure to verify that an authenticator is reachable and responding.

Procedure

1. Make sure the authenticator has been created in Ignition Server as described in [Creating an authenticator](#) on page 104.
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**.

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 97, the *global authenticator* record allows you to create a default RADIUS access policy that applies to requests from unknown authenticators.

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, “globaldefault” and apply policy based on whether the request is being handled by the global authenticator.

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.

Use of a *RADIUS global authenticator* requires an Ignition Base LARGE license.

Creating the Global Authenticator

Follow this procedure to create 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, select the **Accept Requests from Any Authenticator** check box.
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.

RADIUS sub-authenticators

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 *sub-authenticators*. The sub-authenticator 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 sub-authenticators* for the set of logical devices that operate from the IP address. When you create each sub-authenticator, you key it to a RADIUS attribute value. If the RADIUS request contains that value, then the sub-authenticator is used.

Upon receiving a request, Ignition Server finds the authenticator record for the IP address and then chooses the first sub-authenticator whose key value matches a value in the RADIUS request. The sub-authenticator specifies the RADIUS access policy to be used for that logical switch. If no matching sub-authenticator 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 sub-authenticator is used, see [Matching an incoming request to an authenticator record](#) on page 97.

Viewing and editing sub-authenticators

Procedure

1. In Dashboard's **Configuration** tree, expand **Authenticators**. Expand the sub-nodes in the tree until you find the authenticator whose sub-authenticators 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 sub-authenticators of the authenticator.
4. To edit a sub-authenticator, click its name and click **Edit**.

See [Creating a sub-authenticator](#) on page 115 for descriptions of the fields of the Sub Authenticator Details window.

Creating a sub-authenticator

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

Procedure

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 sub-authenticator.

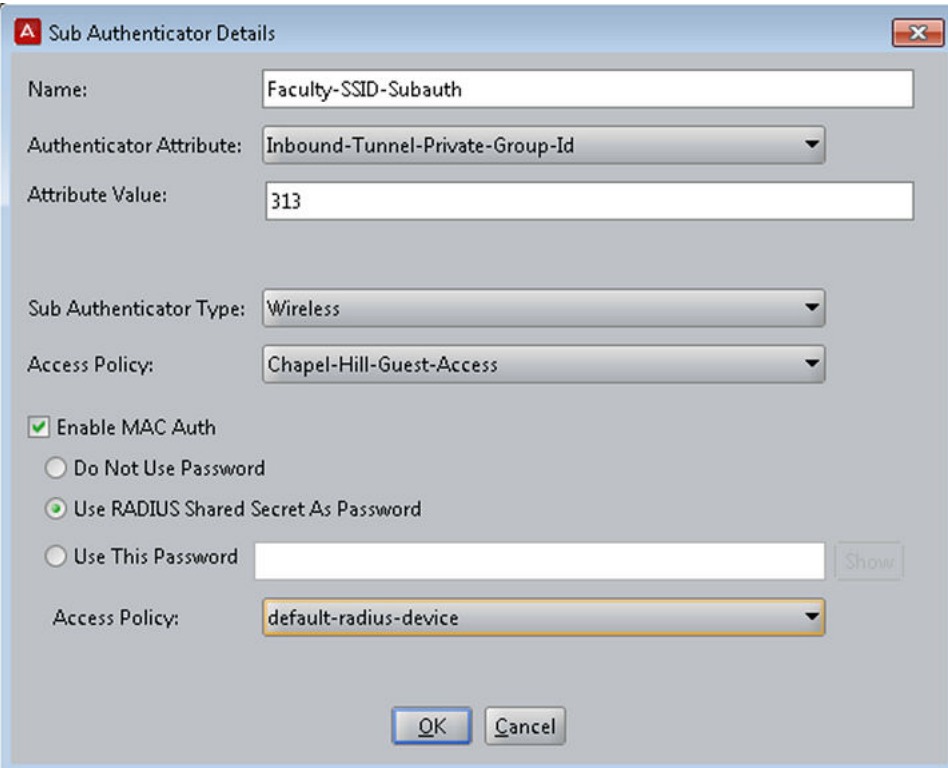
Important:

If you have not defined your authenticator, go to [Creating an authenticator](#) on page 104 and create it now.

2. In the tree, click the authenticator's name.

The **Authenticator Details** panel displays, showing the **Sub Authenticators** panel in its lower half.

3. At the bottom of the panel, click **New**.



Sub Authenticator Details

Name: Faculty-SSID-Subauth

Authenticator Attribute: Inbound-Tunnel-Private-Group-Id

Attribute Value: 313

Sub Authenticator Type: Wireless

Access Policy: Chapel-Hill-Guest-Access

Enable MAC Auth

Do Not Use Password

Use RADIUS Shared Secret As Password

Use This Password Show

Access Policy: default-radius-device

OK Cancel

4. In the **Sub Authenticator Details** window, type a **Name** for the subauthenticator. This name appears in the RADIUS logs and can 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 287.

6. In the **Attribute Value** field, type the RADIUS attribute value that triggers the use of this sub-authenticator.
7. In the **Sub Authenticator Type** drop-down list, choose the type of virtual device that this sub-authenticator represents, such as *Wired*, *Wireless*, or *VPN*.
8. In the **Access Policy** drop-down list, choose the Ignition Server RADIUS access policy that controls user access to this sub-authenticator.
9. If you want to allow MAC Auth on this sub-authenticator, select the **Enable MAC Auth** check box 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 340).
10. Click **OK** to save the sub-authenticator.
11. If you wish to define more sub-authenticators, click **New** again.

Once you have created all the sub-authenticators of this authenticator, you can sort them, as described in [Sorting sub-authenticators](#) on page 116.

Sorting sub-authenticators

When an authenticator has multiple sub-authenticators, 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 sub-authenticator whose attribute/value pair matches a RADIUS attribute/value pair in the request. If any of your sub-authenticators is more widely applicable than others, then you may have to sort the list of sub-authenticators to ensure the desired sub-authenticator takes effect.

Procedure

1. In Dashboard's Configuration tree, expand **Authenticators** and expand the sub-nodes to find the authenticator whose sub-authenticators 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 sub-authenticator 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.

Procedure

1. In Dashboard's Configuration 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 displays.
2. Edit the settings for the selected authenticator.
See [Creating an authenticator](#) on page 104 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

Follow this procedure to disassociate an authenticator from a container.

Procedure

1. In Dashboard's Configuration 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 displays.
2. 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 holds the authenticator.
 - Click **OK**.
3. 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 482.

Follow this procedure to rename an existing authenticator.

Procedure

1. In Dashboard's Configuration hierarchy tree, expand **Authenticators**. Expand the sub-nodes in the tree until you find your authenticator.
2. Click its name and click **Edit**.

The **Authenticator Details** window displays.

3. Enter the new **Name** for the authenticator.
4. 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 482 for troubleshooting instructions.

Follow this procedure to delete an existing authenticator.

Procedure

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

Chapter 8: Using Ignition Access Portal as a Policy Enforcement Point

If you will use Identity Engines Ignition Access Portal as a Policy Enforcement Point for Identity Engines Ignition Server, it is important to understand the specific interactions that occur between the two applications. This section describes those interactions.

For complete information about using Identity Engines Ignition Access Portal, see *Administering Avaya Identity Engines Ignition Access Portal, NN47280-604*.

Inbound and Outbound RADIUS attributes

Inbound and Outbound attributes are key components of any access policy that you configure on the Ignition Server—they determine who is allowed to access a network (authentication) and what they can access upon successful authentication (authorization). It is important to understand which Outbound and Inbound attributes are used for communication between Identity Engines Ignition Access Portal and Ignition Server.

This section describes the attributes that are used between the two applications, and can therefore be included in any Ignition Server access policies that use Identity Engines Ignition Access Portal as an authenticator.

Some attributes have configurable values, while others do not. For information about configuring Identity Engines Ignition Access Portal attributes, see *Administering Avaya Identity Engines Ignition Access Portal, NN47280-604*. Following are the Inbound and Outbound attributes that can be included in Ignition Server access policies for use with Access Portal:

Inbound attributes (from Access Portal to Ignition Server)

RADIUS VSA attributes

- Inbound-Avaya-Access-Portal-Captive-Portal-Zone-Name
- Inbound-Avaya-Access-Portal-IN-Interface-Name

RADIUS standard attributes

- Calling-Station-Id - This attribute contains the MAC address of the authenticating client device.
- Framed-IP-Address - This attribute contains the IP address of the authenticating client device.
- NAS-Identifier - By default, this attribute contains the name of the Access Portal. You can, however, configure the value under Services > Captive Portal > <Zone>, within the RADIUS options section of the Captive Portal(s) tab.
- NAS-Port - This attribute contains a fixed value of 1.
- NAS-Port-Type - This attribute contains a fixed value of 15.

Outbound attributes (from Ignition Server to Access Portal)

Following are the configurable Outbound attributes:

- Outbound-Session-Timeout
- Outbound-Avaya-Access-Portal-Access-Group-Name

Related Links

[Registering Access Portal with the Ignition Server](#) on page 108

[Configuring a guest access policy](#) on page 123

Registering Access Portal with the Ignition Server

Follow this procedure to register Access Portal with the Ignition Server.

Warning:

Any mismatch in RADIUS configuration between the Ignition Server and Access Portal (for example, server IP address, shared secret, password, and so on) can result in fatal or internal errors to the clients. Always perform a test user authentication after configuring RADIUS settings in Ignition Server and Access Portal.

Procedure

1. In the Dashboard **Configuration** tree, expand the **Access Portal** folder and click **Access Portal Servers**.
2. Click **New**.

The **Access Portal Server Details** page displays.

The screenshot shows the 'Access Portal Server Details' configuration window. It includes the following fields and options:

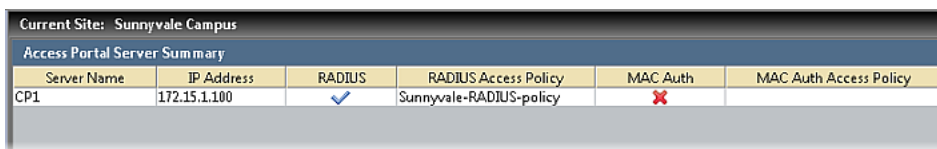
- Name:** Text input field with a red 'x' icon.
- IP Address:** Text input field with a red 'x' icon.
- Trust Device Update**
- Expiration**
 - Date:** 2014-12-28 09:28:19 (with clock and calendar icon)
 - Duration:** Days: 30, Hours: 0
- Delete On Expiry**
- RADIUS Shared Secret:** Text input field with a red 'x' icon and a 'Show' button.
- RADIUS Access Policy:** Dropdown menu showing 'default-radius-user'.
- Enable MAC Auth**
- Access Policy:** Dropdown menu showing 'default-radius-device'.
- Do Not Use Password**
- Use RADIUS Shared Secret As Password**
- Use This Password:** Text input field with a 'Show' button.
- Member Of Groups:** A table with the header 'Internal Group Name' and an empty row below it. Below the table are 'Add...' and 'Remove' buttons.

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. Ensure that you enter the IP address of the ADMIN interface. Also make sure that Access Portal's ADMIN interface is reachable from the Ignition Server.
- **Trust Device Update:** Select this check box if you want the Ignition Server to create a device record in the local store with the device fingerprint of the authenticating user. Note that, 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:** Select this check box if you want to specify an expiry date or lapse period for the devices learned through Access Portal.
 - To specify an expiration date, click **Date** and 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.
 - To specify a lapse period, click **Duration** and use the arrow keys to specify the number of days and hours from the time the device was learned until the time it expires.

- **Delete On Expiry:** Select this check box if you want the Ignition Server to delete learned devices after the expiry date. Note that it may take up to 24 hours after the expiry date for the devices to be purged from the local store.
 - **RADIUS Shared Secret:** Enter the Shared Secret that you configured for RADIUS server.
 - **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-radiususer).
 - **Enable MAC Auth:** Select this check box to provide authentication based on the MAC address of the device that is trying to connect.
 - **Member of Groups:** Select one or more internal groups to which unregistered devices can be auto-associated. Click **Add**, select the internal group or groups, and click **OK**.
4. Click **OK**.

The **Access Portal Server Summary** page displays.



Current Site: Sunnyvale Campus					
Access Portal Server Summary					
Server Name	IP Address	RADIUS	RADIUS Access Policy	MAC Auth	MAC Auth Access Policy
CP1	172.15.1.100	✓	Sunnyvale-RADIUS-policy	✗	

Related Links

[Using Ignition Access Portal as a Policy Enforcement Point](#) on page 120

Configuring a guest access policy

Your guest access policy determines how, when, and where guests can connect to your network, and what sections of your network they can use. If you will use Ignition Guest Manager to create guest user accounts, consult *Avaya Identity Engines Ignition Guest Manager Configuration*, NN47280-501 for instructions.

Use this procedure to create a basic guest access policy.

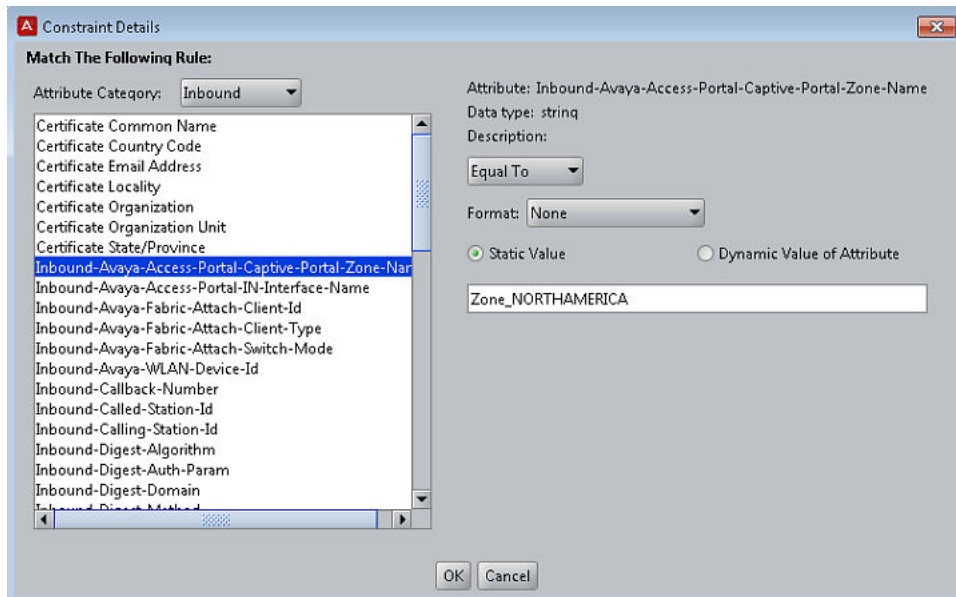
Procedure

1. In the Dashboard **Configuration** tree, click the name of your site.
2. Expand **Site Configuration**, expand **Access Policies** and click **RADIUS**.
3. Click **New**.
4. Enter a name for the new access policy and click **OK**.
5. In the left navigation pane, highlight the name of the new access policy, click the **Authentication Policy** tab and click **Edit**.
6. Configure your tunnel settings. Ensure that you select **PAP** under **None**. Click **OK**.

7. Configure your identity routing policy to enable the Ignition Server to find guest user accounts in the Ignition Server embedded user store. Click the **Identity Routing** tab and click **Edit**.
 - If you already have an identity routing policy that you wish to use, click **Enable Default Directory Set**, and select the **Directory Set** from the drop-down list. Click **OK** to save the policy. Proceed to Step 8.
 - To create a new identity routing policy, do the following:
 - Click **New**.
 - Configure the Ignition Server to use the embedded user store (or any other target directory). In the **Directory Set** section, select **default set** (or any other target set that you wish to use). In the **Match Realm** section, select **Realm Not Specified**. In the **Match Authenticator Container** section, select **Disable Authenticator Container Matching**. Click **OK**.
 - In the **Identity Routing Policy** window, select the **Enable Default Directory Set** check box and select **default set** as the Directory Set. Click **OK**.
8. In the **Access Policy** window, click the **Authorization Policy** tab.
9. In the **RADIUS Authorization Policy** section of the window, click **Edit**.
10. In the **Rules** section, click **Add**.
11. In the **New Rule** window, type a name for the rule and click **OK**.
12. With your rule selected, go to the buttons to the right of the **Constraint** list and click **New**.
13. In the **Attribute Category** drop-down list, select the attribute category **Inbound**.

In response, the list shows all the attributes for **Inbound**.
14. In the list, select one of the following Access Portal Inbound Attributes:
 - RADIUS VSA Attributes**
 - Inbound-Avaya-Access-Portal-Captive-Portal-Zone-Name
 - Inbound-Avaya-Access-Portal-IN-Interface-Name
 - RADIUS Standard Attributes**
 - Calling-Station-Id - This attribute contains the MAC address of the authenticating client device.
 - Framed-IP-Address - This attribute contains the IP address of the authenticating client device.
 - NAS-Identifier - By default, this attribute contains the name of the Access Portal. You can, however, configure the value under Services > Captive Portal > <Zone>, within the RADIUS options section of the Captive Portal(s) tab.
 - NAS-Port - This attribute contains a fixed value of 1.
 - NAS-Port-Type - This attribute contains a fixed value of 15.
15. Select the appropriate value options and enter the value for the selected attribute.

In this example, the Inbound attribute “Inbound-Avaya-Access-Portal-Captive-Portal-Zone-Name” is used with a value of “Zone_NORTHAMERICA”. Zone_NORTHAMERICA has two IN interfaces associated with it. This rule therefore applies to all users who enter through either of the two IN interfaces associated with Zone_NORTHAMERICA.



16. Click **OK** to close the Constraint Details window and return to the Edit Authorization Policy window.

17. In the **Action** section of the Edit Authorization Policy window, click **Allow**.

A list of available Outbound Values displays. Any Access Portal Access Groups that have been created are listed as available Outbound Values.

18. Select one of the following Access Portal Outbound Values using the arrows to move the desired values into the **Provision With** field:

- <Access Group Name>
- Session-Timeout

19. Click **OK**.

In this example, “Access-Group-Guest” is the selected Outbound Value. That is, all users who enter through either of the two IN interfaces associated with Zone_NORTHAMERICA will be granted access through the OUT interface, and see the success page associated with the Access Portal Access Group named “Access-Group-Guest”.

Using Ignition Access Portal as a Policy Enforcement Point

Selected Rule Details

Rule Name: Rule Enabled

(Constraint)	AND/OR
▼	Inbound.Inbound-Avaya-Access-Portal-Captive-Portal-Zone-Name = Zone_NORTHAM...	▼	▼

Action

Allow
 Deny
 Check Posture
 NAP

Provisioning (Outbound Values)

Provision With

All Outbound Values

- Access-Group-Name-BusinessPartner
- Access-Group-Name-Employee
- Access-Group-Name-Guest
- Access-Group-Name-SantaClara-Em...
- Admin-Access
- Guest-Man

Summary

IF Inbound.Inbound-Avaya-Access-Portal-Captive-Portal-Zone-Name = Zone_NORTHAMERICA
THEN **Allow**
Send Outbound Values: Access-Group-Guest

Related Links

[Using Ignition Access Portal as a Policy Enforcement Point](#) on page 120

Chapter 9: 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.

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 148), 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.

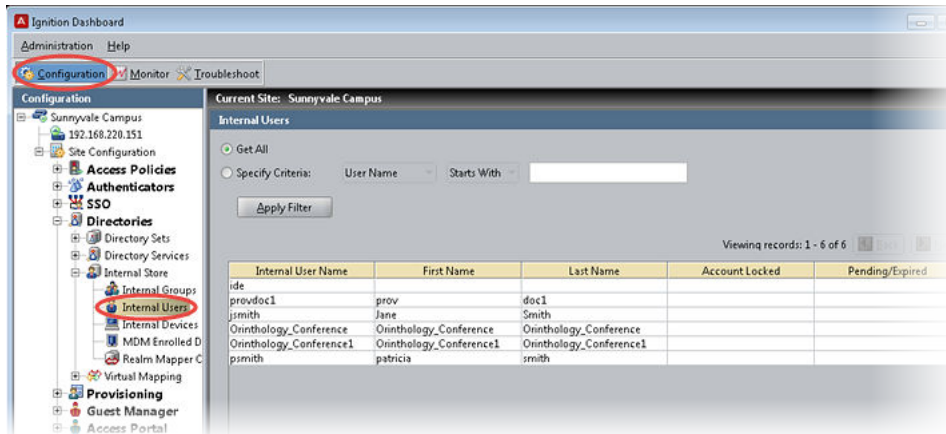
Internal Users Panel

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 to:

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

! Important:

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.



From this panel you can.

- View the list of all users in the internal store.
- Add, edit, copy or delete internal users.

To do this, click the appropriate command button at the bottom of the panel.

- Sort the list of users by a particular column.

To do this, 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 129.

- 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.

Procedure

1. In the **Internal Users** window, select the **Specify Criteria** check box.
2. Two drop-down lists are shown to the right of the **Specify Criteria** check box. 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

Follow this procedure to view the complete details of an internal user.

Procedure

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.
3. Click **Edit** or double-click on the desired user entry in the displayed list.

Ignition Dashboard displays the **Edit User** window for the selected user. The **Edit User** window shows all the data for a selected user.

4. 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](#) on page 129.

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.

Follow this procedure to create a new internal user.

Procedure

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 User** window displays.

- **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 to. Then, click **New** in the **Users** tab of the **Internal Group Details** panel. The **Edit User** window displays.

The screenshot shows a 'New Internal User' dialog box with the following fields and settings:

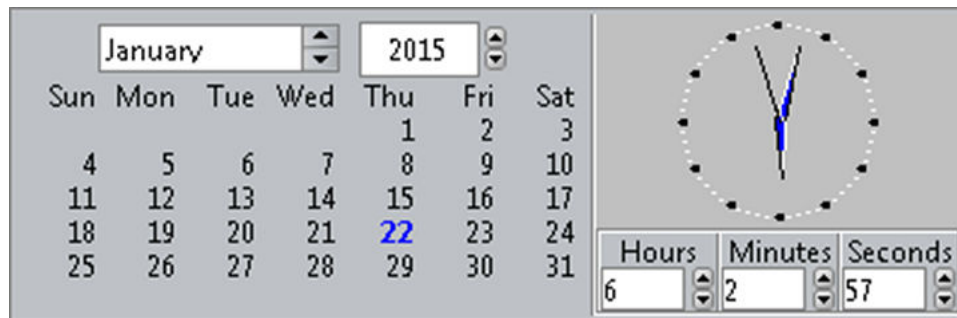
- Info:**
 - User Name: sclemens
 - First Name: Samuel
 - Last Name: [empty]
 - Password: [masked]
 - Confirm Password: [masked]
 - Start Time: 2015-01-14 22:58:24
 - Password Expires: 2016-01-14 22:58:24
 - Max Retries: 3
 - Account Locked:
 - Delete on Expire:
- Custom Attributes:**
 - Title: [empty]
 - Network Usage: [empty]
 - Email Address: sclemens@company.com
 - IPv4 Address: [empty]
 - Org. Role: [empty]
 - Office Location: Sunnyvale
 - Comments: [empty]
- Member Of Groups:**
 - Internal Group Name: [empty list]
 - Buttons: Add..., Remove...

2. In the **Edit** window, enter the user details and select the appropriate check box settings for the new user account. The fields and settings that describe the user are as follows.

- **User Name, First Name, Last Name:** The login name, given name, and family name of the user, respectively.
- **Account Disabled:** Select this check box 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 check box 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.



- **Enable Password Expire:** Select this check box 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.
- **Delete on Expire:** If you want to have Ignition Server automatically delete the account after it expires, select this check box. 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 check box 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 256.

*** Note:**

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 144. 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 358.

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 135.

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

Importing user records

Procedure

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

```
User Name,First Name,Last Name>Password,Start Time>Password Expires,
Max Retries>Delete on Expire,Account Disabled>Title,Role,
Network Usage,Office Location>Email Address,Comments,
Group Name,Devices,Custom ACLs,Custom VLANS,Provisioners Groups,
isGuest,isProvisioner,enableStartTime,enablePwdExpTime,
enableMaxRetries,isSelfProv,activateOnFirstLogin,Credential Type,
Provisioned By,Prov Domain,Prov Group,Custom IP,Custom Number,
Creation Time,Sponsor First Name,Sponsor Last Name,Sponsor Email Address,
Sponsor Cellphone,Sponsor Response,Guest Details
```

Important:

The fields Creation Time, Sponsor First Name, Sponsor Last Name, Sponsor Email Address, Sponsor Cellphone, Sponsor Response, and Guest Details are added in Release 9.1. You can successfully import Release 9.0 exported user records to Release 9.1. Make sure that the groups already exist in Ignition Server and the newly added fields will show the default values.

2. In the Dashboard's Configuration hierarchy tree, expand **Directories**, expand **Internal Store**, and click **Internal Users**.
3. On the Internal Users page, click **Import**.
4. In the User Record Import window, click **Browse**, navigate to find your CSV file, select it, and click **Open**.
5. Select **Override duplicate records** to override user records with the same User Name.
6. Click **OK** to import the user records.

The Internal Users page displays the imported users.

Exporting user records

Procedure

1. In the Dashboard's Configuration hierarchy tree, expand **Directories**, expand **Internal Store**, and click **Internal Users**.
2. On the Internal Users page, click **Export**.
3. In the Export User Records window, do one of the following:
 - To export all user records, select **Get All**.
 - To select some user records, select **Specify Criteria** and set your filter criteria in the fields to the right.

In the first drop-down list, select the name of the attribute that 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 the records whose attribute is identical to the whole search string. Type the search string in the field at the right.
4. Click **Browse**, navigate to 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**.
5. In the Export User Records window, click **OK** to export the records.

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 [Introduction to MAC Authentication](#) on page 339); 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 [Introduction to Asset Correlation](#) on page 353).

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 [Windows Machine authentication](#) on page 316.

! **Important:**

If an internal device becomes learned from a Mobile Device Management server, Ignition server removes the device record from the Internal Devices list as the device is included in the MDM Enrolled Devices. For more information on MDM enrolled devices, see [MDM enrolled devices](#) on page 196.

Finding an internal device

Follow this procedure to find a device record.

Procedure

1. In Dashboard's Configuration tree, click your site, expand **Site Configuration > Directories > Internal Store**, and click **Internal Devices**.

The Internal Devices panel shows your device records. See "[Filtering the device list](#)" on page 134 for instructions on finding a device in the list.

2. Use the **Back** and **Next** buttons to move through the list.

MAC Addresses are stored only in the internal store

In Ignition Server, 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

Follow this procedure to filter the Internal Devices panel.

Procedure

1. In the **Internal Devices** panel window, click **Specify Criteria**.
2. Two drop-down lists are shown to the right of the Specify Criteria check box. 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. 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 Server. (If you need to create *many* device definitions, you may prefer to create them in bulk as shown in [Importing device records](#) on page 138.)

Procedure

1. In Dashboard's Configuration tree, click on the name of your site, expand **Site Configuration > Directories > 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 can write the twelve characters without separators, or you can 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, select the **Record Disabled** check box.
4. In the **Name** field, type a name for the device.

This name identifies the device in logs and when you associate it with a group or user.

5. If you want Ignition Server to delete this record automatically after its expiration date, select the **Delete on Expire** check box.

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, and so on.

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 138). 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 to 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** check box 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.

15. The **Custom Attributes** fields allow you to record additional information about the device.

See [Adding Virtual Attributes for Devices](#) on page 235.

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 using MAC authentication (bypassing 802.1X), then make sure you have a MAC authorization policy that applies to it. See [MAC authentication set-up procedure example](#) on page 343.
- If this device is to be assigned to a user in order to enforce an asset correlation policy, see [Assigning a device to a user or group](#) on page 137.

Assigning a device to a user or group

You can 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 135.

Procedure

1. In Dashboard's Configuration tree, click on the name of your site, expand **Site Configuration** > **Directories** > **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. Alternatively, click it and click **Edit**).
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, select the check box 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 354.

* **Note:**

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: 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 tree, click on the name of your site, expand **Site Configuration > Directories > 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 135.

Deleting a device record

To delete a device record, in Dashboard's Configuration tree, click on the name of your site, expand **Site Configuration > Directories > 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

Follow this procedure to import device records.

Procedure

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,Attribute 1,Attribute 2,Attribute 3,
Attribute 4,Attribute 5,Attribute 6,"Group1,Group2","User1,User2",
Account Disabled,Start Time,End Time,isGuest?,ActivateOnFirstLogin?,
DeleteOnExpire?,EnableStartTime?,EnableExpTime?,Provisioned By,
Prov Domain,Prov Group
```

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 attribute values, membership in the groups *default* and *printers-in-HQ*, and the **Record Disabled** check box is not selected. Make sure the groups exist already in Ignition Server.

! Important:

The fields Provisioned By, Prov Domain, and Prov Group are added in Release 9.1. You can successfully import Release 9.0 exported device records to Release 9.1. Make sure that the groups already exist in Ignition Server and the newly added fields will show the default values.

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 **Browse**, 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. Edit the name if desired. This name is saved as the Source attribute in each device record.
6. Select **Override duplicate records** to override device records with the same Mac Address.
7. Click **OK** to import the records.

The Internal Devices page displays the imported devices.

Exporting device records

To export a set of device records, use the following procedure.

Procedure

1. In Dashboard's Configuration tree, click on the name of your site, expand **Site Configuration > Directories > Internal Store**, and click **Internal Devices**.
2. In the Internal Devices panel, click **Export**.
3. In the **Device Record Export** window, do one of the following.
 - To export all device records, select **Get All**.
 - To export some device records, select **Specify Criteria** and set your filter criteria in the fields 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.

4. Click **Browse**, 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**.

5. In the **Device Record Export** window, click **OK** to export the records.

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.

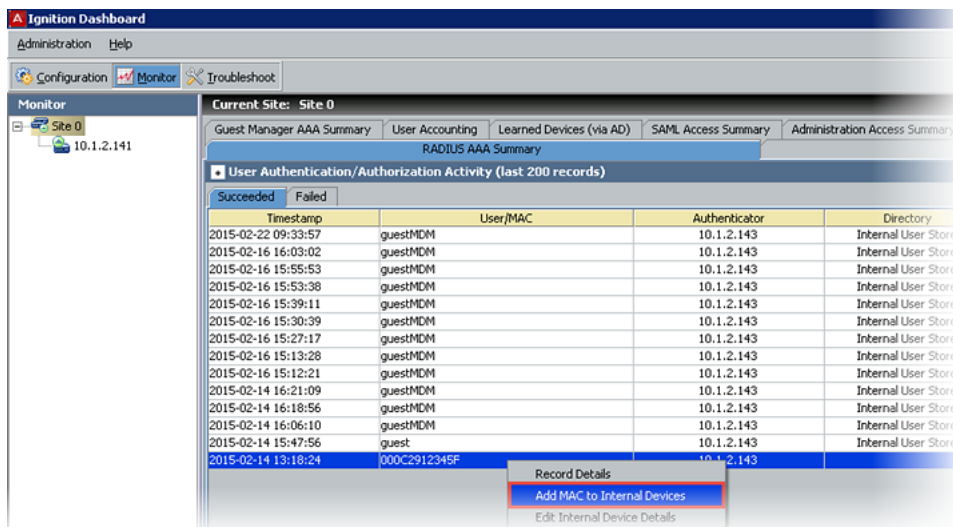
To do this, use the MAC address wildcarding feature, define your device as usual in the Device Record window, but specify a partial MAC address followed by a "*" character.

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 356 for details on asset correlation policies.

The screenshot shows a window titled "b9:4a* - Device Record Details". Under the "Info" section, there are several fields: "MAC Address" (containing "b9:4a*", circled in red), "Name" (containing "AI- Xerox-printers"), "Type" (a dropdown menu set to "printer"), "Source" (an empty text box), "VLAN Label" (containing "hq-printer-vlan"), and "VLAN ID" (containing "206"). To the right of the "MAC Address" field is a checkbox labeled "Record Disabled" which is currently unchecked. Below the "Info" section is a section for "Custom Attributes".

Adding an internal device from the Monitoring Access Logs Procedure

1. In the Monitor tab, right-click on a log record from the **Succeeded** or **Failed** tab.



2. Select **Add MAC to Internal Devices**.

The MAC address is automatically added and the device details can be edited in the **New Device Record** window.

3. Click **OK** to complete adding the new device record.

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.

Internal groups panel

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

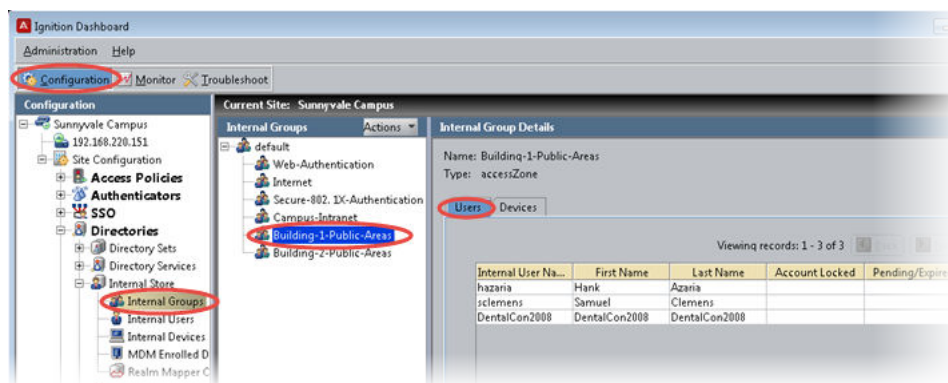
This window consists of

- **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.



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

! Important:

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, the Ignition Dashboard does not permit this default group to be renamed or to be deleted.

Adding a new internal group

Follow this procedure to add a new internal group.

Procedure

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 Guest Manager Configuration*, NN47280-501 for more information.

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. Follow this procedure to move an internal group.

Procedure

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

Follow this procedure to rename an existing internal group.

Procedure

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 and 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 Guest Manager Configuration*, NN47280-501 for details.

Follow this procedure to change the group type.

Procedure

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.

Follow this procedure to delete an existing group.

Procedure

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 selected check box 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 selected check box 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. Remove all users from the Group. Click each user row and click **Delete**.
 - Click the Devices tab. Remove all devices from the Group. Click each user row and click **Delete**.
5. To delete the group, click **Actions > Delete Internal Group...**
6. Click **Yes** to delete the selected group.

For more information about virtual groups, see [Virtual Groups and Attributes](#) on page 223.

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.

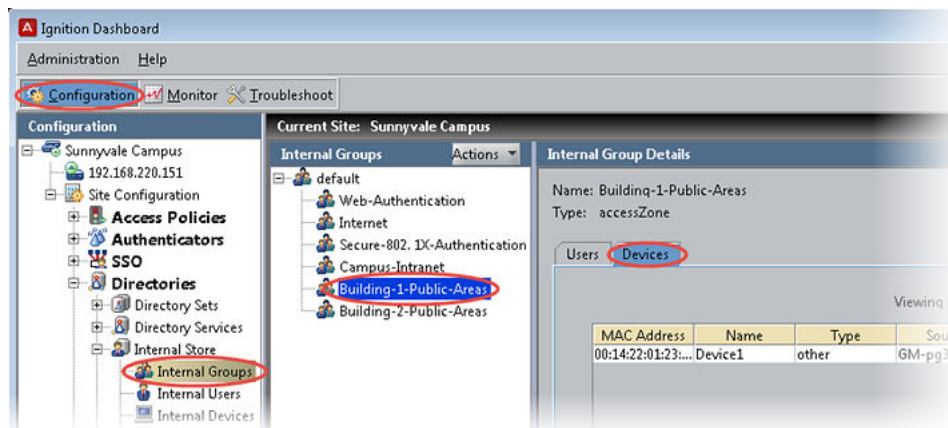
Users tab

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. The following example shows 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.

Chapter 10: 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.

Quickstart: Directory Services in Dashboard

Three tabs in Dashboard allow you to perform operations 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, and 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 150.

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 186.

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 186.

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 more information, see [How Ignition Server looks up a user for Authentication and Authorization](#) on page 246.

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 Server. In Ignition Server, 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 configured directory services, authentication servers (if necessary), and directory sets, you create identity routing policies as explained in [Understanding Identity Routing Policy](#) on page 246.

Directory Services

A directory service establishes Ignition’s connection to a user repository or authentication server. The directory service can 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 Services](#) on page 148.)

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
- Citrix XenMobile MDM server

For a list of which authentication protocols are supported using which directory servers, see [Managing Ignition Server licenses](#) on page 75.

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.

Important:

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 [Internal users, groups, and devices](#) on page 127.

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 148, 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. The commands are explained in [Checking directory service connections](#) on page 186.

Troubleshoot View Dashboard's **Troubleshoot** view lets you test user authentications and lookups. See [Troubleshooting user lookup and authentication](#) on page 186 .

Connecting to active directory

This section explains how to connect to Active Directory.

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.

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	_____
<p>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 <i>Ignition Server service account</i>. If you want to perform MSCHAPv2 authentication, the service account must have permission to <u>create</u> and <u>delete</u> computer accounts (the <i>Create Computer Object</i> and <i>Delete Computer Object</i> permissions) in the <i>Netlogon account root</i> in Active Directory. See “Netlogon Account Root DN” in this table. If you have not specified a Netlogon account root DN in Ignition, then the service account must have these permissions in the <i>Computers container</i> of your AD service.</p> <p>Ignition Server uses the service account to join the Active Directory domain. Joining the domain requires creating a machine account in the <i>Netlogon account root</i> 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.</p> <p>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 <i>COMPANY.COM/Administrator</i> is usually <i>Administrator</i>.</p> <p>Creating a service account: If no appropriate account exists in your AD installation, see “Create the Service Account in AD” in the <i>Ignition Server Getting Started Guide</i>. For help setting its permissions, see “Set the</p>	

Table continues...


Setting Name	Setting Value
AD Permissions of the Service Account” in <i>Avaya Identity Engines Ignition Server Getting Started Configuration, NN47280–300</i> .	
Service Account Password	
The Service Account Password is the password for the AD service account. <i>Do not record the password here.</i>	
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)	
The LDAP Port of the primary and secondary AD data stores. For SSL, enter 636. If SSL is not used, enter 389. You <i>cannot</i> use the global catalog port (3268). <i>Use the LDAP ports (389 and 636) only!</i>	
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 286.	
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 <i>Active Directory</i> stores. For instructions on using Microsoft tools to find this name, see Looking up AD settings: Finding Domain and NetBIOS names on page 160.	
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 159 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 159.	

Table continues...

Setting Name	Setting Value
Netlogon Account Root DN	_____
<p>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.</p> <p>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 <i>Avaya Identity Engines Ignition Server Getting Started Configuration, NN47280–300</i></p>	

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 151.
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).

Warning:

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 NTLMv1 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 *Avaya Identity Engines Ignition Server Getting Started Configuration, NN47280–300*.
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 *Avaya Identity Engines Ignition Server Getting Started Configuration, NN47280–300*.
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 procedure 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 157.

Before you begin

- Make sure your DNS server addresses, clock setting, and firewall settings have been completed as explained in [Preparing to connect to Active Directory](#) on page 153.
- Make sure your primary directory server is reachable. The wizard connects to it in order to retrieve group and schema information.

Procedure

1. In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**. Click **New**.

Ignition Server launches the **Create Directory Service Wizard**. The **Choose Service Type** window displays.

2. Select **Active Directory** and click **Next**.

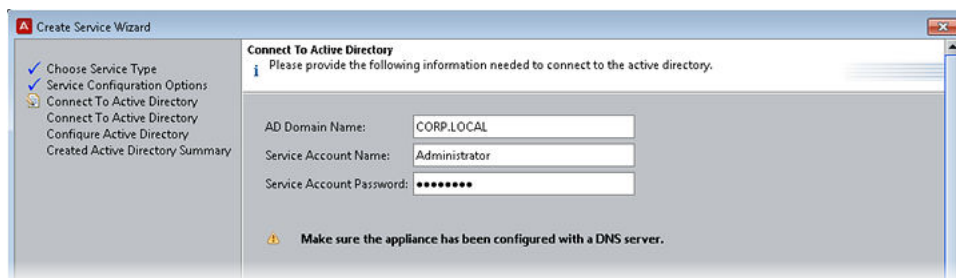
The **Active Directory Configuration Options** window displays.

3. Select **Automatically Configure** and 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.

After you type these settings, click **Next**.



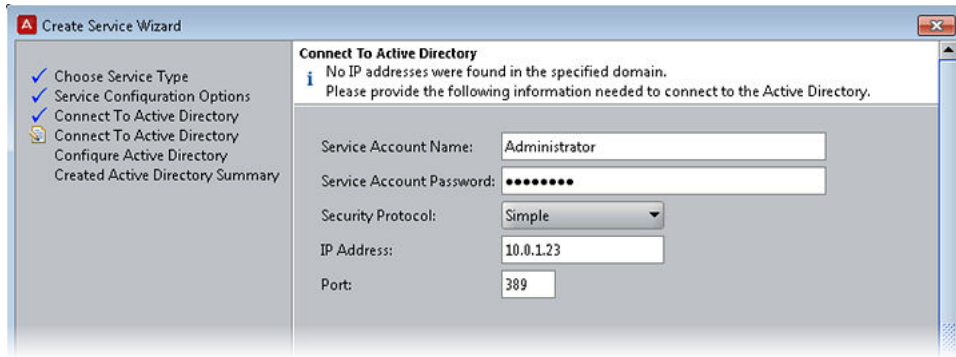
5. In the next window:

- Choose the **Security Protocol**.

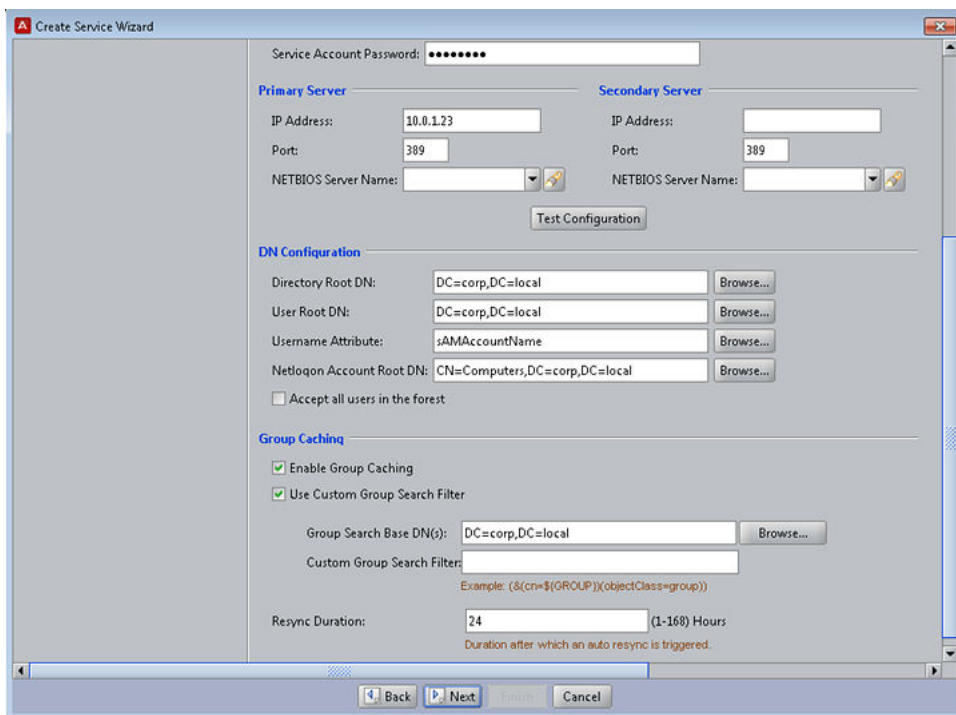
Choose **Simple** for unencrypted communication with AD or choose **SSL** for encrypted communication.

- 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).



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



6. In this window:

- **Name:** Assign the directory service a name in the **Name** field.
- If needed, you can 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 151.
- **Primary Server:** In the **Primary Server** section, specify the **NETBIOS Server Name**.
- **Secondary Server:** Add a **Secondary Server** if desired. This is a backup AD server.
- **DN Configuration :** 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 can 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.

- **Netlogon Account Root DN:** 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 [AD connection settings](#) on page 151.
- **Enable Group Caching:** The Ignition Server maintains an internal cache of the group hierarchies and attributes schemas of the directory services. If necessary, disable this caching by clearing the **Enable Group Caching** check box.
- **Group Search Base DNs:** By default, Ignition Server looks for groups beginning at the Directory Root DN. You can change this default behavior by specifying **Group Search Base DNs**. This is useful in case of huge AD deployments, where beginning at the root DN can take a substantial amount of time. In addition, you can restrict the types of groups that IDE caches by specifying a custom Group Search Filter. The filter follows the LDAP query syntax.
- **Re-sync Duration:** Enter the sync interval between Ignition Server and Active Directory, in hours, in **Re-sync Duration**. The range is 1 to 168 hours. The cache is automatically refreshed based on this setting.

7. Click **Test Connections**.

Testing the connection can 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 486.

8. Click **Next**.

The next window summarizes the connection settings of the service.

9. Click **Finish**.

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

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 154.

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 68.
2. Make the clock and firewall settings explained in [Connecting to active directory](#) on page 151.

3. Make sure your primary directory server is reachable. The wizard connects to it in order to retrieve group and schema information.

Follow this procedure to connect to AD (manual mode).

Procedure

1. In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**. Click **New**.

Ignition Server launches the Create Directory Service Wizard. The **Choose Directory Service Type** window appears.

2. Select **Active Directory** and click **Next**.

The **Active Directory Configuration Options** window displays.

3. Select **Manually Configure** and click **Next**.

The **Configure Active Directory** window displays.

4. 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 151.

5. Click **Test Connections**.

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 486.

6. Click **Next**.

A window appears, summarizing the settings you have made.

7. 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 186 for an explanation of the icons.

Troubleshooting AD connections

Diagnosing connection problems

Check your AD connection.

Procedure

1. Use the **Test Connections** or **Recheck Service** button. See [Troubleshooting user lookup and authentication](#) on page 186.

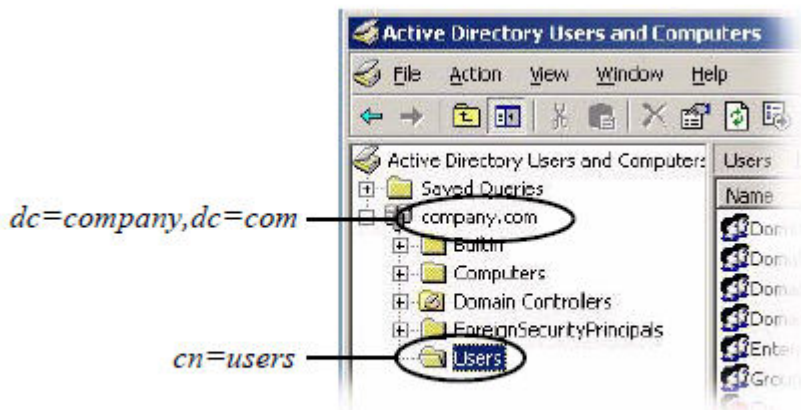
2. Check your AD settings (listed in [AD connection settings](#) on page 151).
3. Check your directory service connection using the Advanced Troubleshooting window.
 - Place the directory service in a directory set (See [Adding directories and authentication servers to a directory set](#) on page 183.)
 - Use the **Process Request** tab and Test Join feature of the Directory Service debugger. See [Troubleshooting user lookup and authentication](#) on page 186.)

Looking up AD settings: Finding your Root DNs

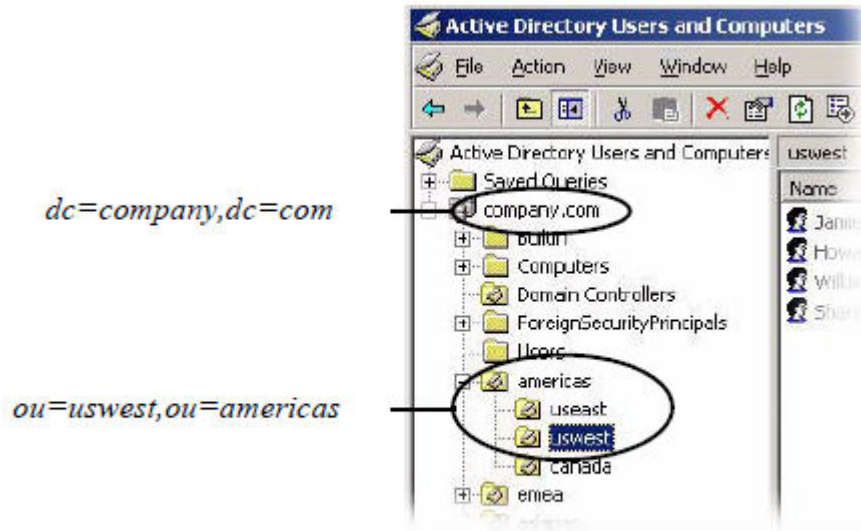
User Root DN and **Directory Root DN**: Enter the names of containers in your AD data store using X.500 naming. **User Root DN** points to the AD container that stores your user records. **Directory Root DN** points to the root of your AD tree and is used to obtain schema and group information.

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.

In this example, User Root DN is cn=users,dc=company,dc=com.



In this example, User Root DN is ou=uswest,ou=americas,dc=company,dc=com.



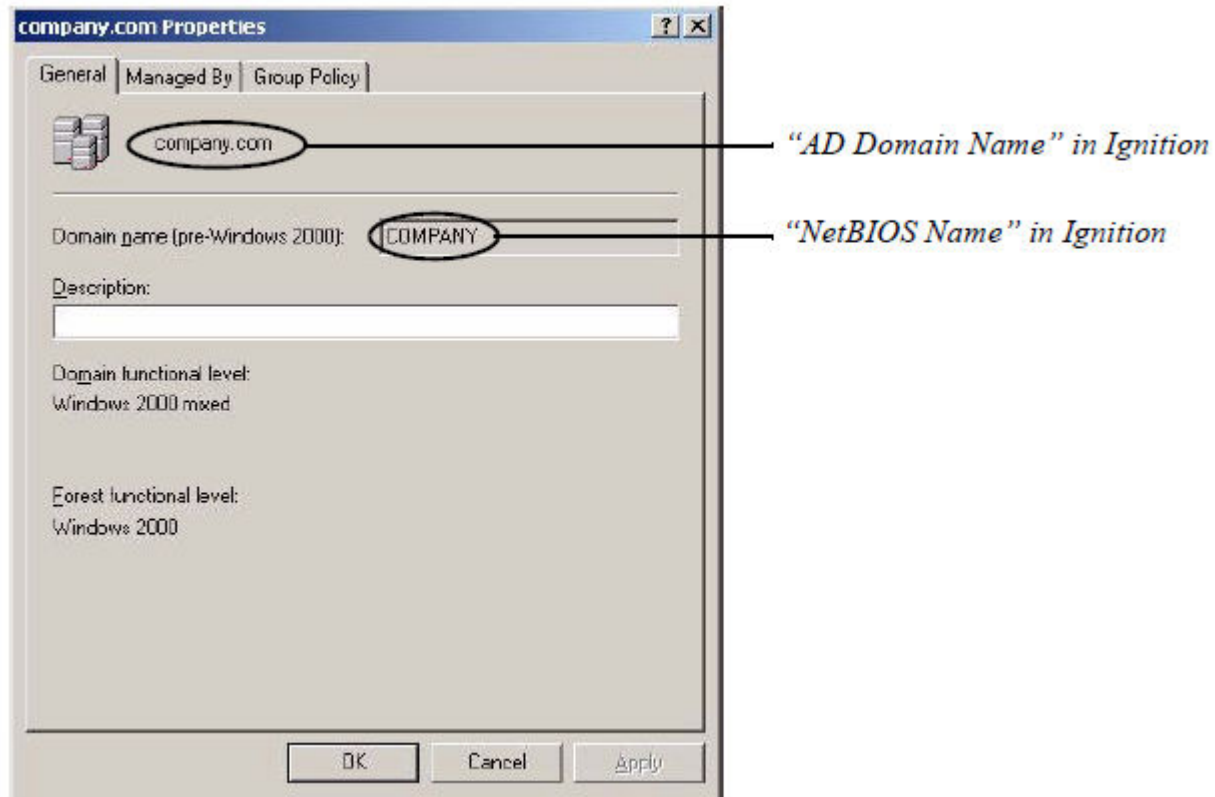
You form the full User Root DN name by adding the CN or OU portion of the name as a prefix to the root portion of the name, as shown in the two examples above. The following example text uses “cn=users,dc=company,dc=com” as our example DN.

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 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 478.

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 149.

LDAP connection settings

The following table describes the parameters Ignition Server uses to connect to an LDAP service. Configure these settings in the Create Directory Service Wizard or in the **Directory Server Details** window. The following table contains an 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, <code>dc=starironinc,dc=com</code> .
Browse Buttons	The Directory Root DN, User Root DN, and Username Attribute buttons allow you to browse your schema to set those values. <div style="display: flex; align-items: flex-start;"> <div style="margin-right: 10px;"></div> <div> <p>Note:</p> <p>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.</p> </div> </div>
Directory Root DN	DN where the LDAP schema containing your users and groups are found. For example, <code>dc=company,dc=com</code> . When you connect the directory service, the Ignition Server Create Service wizard attempts to choose a Directory Root DN for
MSCHAPv2 authentication	LDAP only: check box indicating whether this LDAP store supports MSCHAPv2 authentication. Also see LDAP Password Attribute in this table.
Name	Each directory service you 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 a policy that provides special provisioning for users who authenticate against a particular directory. See Inbound Attributes on page 286.

Table continues...


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 170.
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). 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. <ul style="list-style-type: none">• 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, <code>cn=Directory Manager</code>
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 <code>jsmith@company.com</code> would be submitted to LDAP as <code>jsmith</code> .

Table continues...

Field Name	Description
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.

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 procedure shows how to use the *automatic mode* of the wizard, which retrieves the default DNSs for you. For instructions on using the manual mode, see [Creating an LDAP Directory Service: Manually configuring](#) on page 167.

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 68.
2. Make sure your primary directory server is reachable. The wizard connects to it in order to retrieve group and schema information.

Procedure

1. In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, click **Directory Services**, and click **New**.
2. Ignition Server launches the **Create Directory Service Wizard**.
The **Choose Directory Service Type** window displays.
3. Select the desired type of LDAP server and click **Next**.
4. Click **Automatically Configure** and click **Next**.
5. In the **Connect To...** window, enter your LDAP administrator credentials in the Service **Account DN** and **Service Account Password** fields.

Type the IP address and port of the LDAP directory. If you want to establish an encrypted connection to LDAP, select **Use SSL**. Click **Next**.

Create Directory Service Wizard

Connect To Sun Directory Server

Please provide the following information needed to connect to the Sun Directory Server.

Service Account DN:

Service Account Password:

Use SSL: Use SSL

IP Address:

Port:

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

Create Service Wizard

Configure Generic LDAP

Created Directory Service Summary

Settings

Name:

Service Type: Generic LDAP

Use SSL: Use SSL

Service Account DN:

Service Account Password:

Directory Root DN:

User Root DN:

Username Attribute:

Use User Search Filter

Example: (&(objectclass=person)(uid=\$USER))

MSCHAPv2 Authentication

LDAP Password Attribute:

Strip Realm

Primary Server **Secondary Server**

IP Address: IP Address:

Port: Port:

Group Caching

Enable Group Caching

Use Custom Group Search Filter

Group Search Base DN(s):

Custom Group Search Filter:

Example: (&(cn=\$(GROUP))(objectClass=group))

Resync Duration: (1-168) Hours

Duration after which an auto resync is triggered.

6. In this window, edit fields as needed.

For an explanation of each field, see [LDAP Connection Settings](#) on page 162.

- 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.

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, select the **Strip Realm** check box.

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), select the **MSCHAPv2 Authentication** check box and click **Browse** to select the name of the LDAP user attribute that holds the NT-hashed password, or type the attribute name directly in the **LDAP Attribute** field.

See [Setting up MSCHAPv2 authentication on LDAP](#) on page 170.

- If needed, edit the **Primary/Secondary Server** settings.
- The Ignition Server maintains an internal cache of the group hierarchies and attributes schemas of the directory services. If necessary, disable this caching by clearing the **Enable Group Caching** check box.
- By default, Ignition Server looks for groups starting at the Directory Root DN. You can change this default behavior by specifying **Group Search Base DNs**. This is useful in case of huge AD deployments, where beginning at the root DN can take a substantial amount of time. In addition, you can restrict the types of groups that IDE caches by specifying a custom Group Search Filter. The filter follows the LDAP query syntax.
- In **Re-sync Duration**, enter the sync interval between Ignition Server and Active Directory in hours.

The range is 1 to 168 hours. The cache is automatically refreshed based on this setting.

7. Click **Test Configuration**.

Testing the connection can 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 486.

8. Click **Next**.

The next window summarizes the connection settings of the service.

9. Review the settings and if correct, click **Finish**.

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

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 164.

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 68.
2. Make sure your primary directory server is reachable. The wizard connects to it in order to retrieve group and schema information.

Procedure

1. In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**. Click **New**.

Ignition Server launches the Create Directory Service Wizard. The **Choose Directory Service Type** window displays.

2. Select the desired type of LDAP server. Click **Next**.
3. Click **Manually Configure**. Click **Next**.
4. In the Connect to LDAP window (specific to the type of LDAP store that you selected), do the following:

Connect To Generic LDAP

i Please provide the following information needed to connect to the Generic LDAP.

- a. In the **Service Account DN** field, enter the DN of the LDAP administrator account. Ignition Server will connect as this administrator. For example, cn=Directory Manager.
- b. In the **Service Account Password** field, enter the password of the LDAP administrator.
- c. **Use SSL:** 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 will travel over the network in unencrypted form. Avaya strongly recommends using an SSL connection to connect to your directory server.
- d. In the **IP Address** field, enter the IP address of the primary LDAP server.
- e. In the **Port** field, enter the Port number at which the LDAP service can be reached. When Use SSL is selected, the Port Entry is typically 636. When Use SSL is not selected, the Port Entry is typically 389.

5. Click **Next**.

The Configure LADP window appears.

6. In the Settings section, edit the fields as needed. For an explanation of each field, see [LDAP Connection Settings](#) on page 162.

- 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 **Browse** and selecting the proper root.

! Important:

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, select the **Strip Realm** check box. 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), select the **MSCHAPv2 Authentication** check box and click **Browse** 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 170.
7. The **Primary Server IP Address** and **Port** fields are populated by the wizard; if necessary, click the padlock button to unlock and then click in the fields to edit them.

The **Secondary Server IP Address** and **Port** fields are optional. If you have a backup server, enter its address here.

The screenshot shows two columns of configuration fields. The left column is titled 'Primary Server' and contains two text boxes: 'IP Address:' with the value '10.177.211.128' and 'Port:' with the value '389'. The right column is titled 'Secondary Server' and contains two text boxes: 'IP Address:' which is empty and 'Port:' with the value '389'. Below these fields is a button labeled 'Test Configuration'.

8. In the Group Caching section
- a. The Ignition Server maintains an internal cache of the group hierarchies and attribute schemas of the directory services. If necessary, disable this caching by clearing the **Enable Group Caching** check box.
 - b. By default, Ignition Server looks for groups starting at the Directory Root DN. You can change this default behavior by specifying **Group Search Base DNs**. This is useful in case of huge deployments, where starting at the root DN can take up a substantial amount of time. In addition, you can restrict the types of groups that IDE caches by specifying a custom Group Search Filter. The filter follows the LDAP query syntax.
 - c. Enter the sync interval between Ignition Server and the LDAP service, in hours, in **Resync Duration**.

The range is 1 to 168 hours. The cache is automatically refreshed based on this setting.

The screenshot shows the 'Group Caching' configuration section. It includes a checked checkbox for 'Enable Group Caching' and an unchecked checkbox for 'Use Custom Group Search Filter'. Below these is a text box for 'Group Search Base DN(s):' containing 'dc=genetics,dc=wustl,dc=edu' and a 'Browse...' button. There is also a text box for 'Custom Group Search Filter:' which is empty. Below that is an example filter: '(?(cn=*)(GROUP)(objectClass=group))'. At the bottom, there is a text box for 'Resync Duration:' containing '24' and '(1-168) Hours'. A note below the text box reads: 'Duration after which an auto resync is triggered.'

9. Click **Next**.

The wizard applies your settings to create the directory service in Ignition Server and displays the confirmation page.

10. Review the settings. If the settings are correct, click **Finish** to create the directory service.

Your directory service has been saved in Ignition Server.

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 186 for an explanation of the icons.

Creating an MDM directory service

For information on creating an MDM directory service, see [Creating an MDM directory service](#) on page 193.

Creating a Token authentication service

For information on connecting to hardware token authentication service, see [Overview of Token Authentication in Ignition](#) on page 202.

Creating a Kerberos Authentication service

For information on connecting to a Kerberos server, see [Setting up a Kerberos Authentication Service](#) on page 201.

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 through MSCHAPv2 to a non-AD LDAP store. This feature is called “MSCHAPv2 termination against LDAP.”

The following sections explain how to configure 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 240.

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

- [MSCHAPv2 termination using an LDAP Password Attribute](#) on page 171
- [MSCHAPv2 termination using a Novell universal password](#) on page 174
- [MSCHAPv2 termination using an OID authentication descriptor](#) on page 175

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

To configure MSCHAPv2 authentication against an LDAP directory service, do the following.

Procedure

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 173.
3. In Ignition, create or edit your LDAP directory service and make these settings in the Directory Server Details window.
 - Select the **MSCHAPv2 Authentication** check box.
 - If necessary, select the **With LDAP Password Attribute** check box. (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 **Browse** and select the attribute,

The screenshot shows the 'edir-mv1 - Novell eDirectory Details' configuration window. The 'Settings' section includes the following fields and options:

- Name: edir-mv1
- Service Type: Novell eDirectory
- Security Protocol: Use SSL
- Service Account DN: cn=admin,o=school.edu
- Service Account Password: [masked]
- Directory Root DN: [empty] Browse...
- User Root DN: [empty] Browse...
- Username Attribute: uid Browse...
- Strip Realm
- MSCHAPv2 Authentication
 - With LDAP Password Attribute: password-mschapv2 Browse...
 - With Universal Password

At the bottom, there are sections for 'Primary Server' and 'Secondary Server'.

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.

A given user might have multiple protocols available to him or her for logging in. For example, you can configure 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 240.

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 following procedure shows you how to deploy mechanisms that create and maintain the NT hash.

Procedure

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. Configure 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. **Configure 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 through MSCHAPv2 against LDAP.

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.

To configure MSCHAPv2 authentication against Novell eDirectory using the eDirectory Universal Password feature, use the following procedure.

Procedure

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 [Creating an LDAP Directory Service Automatically Configuring](#) on page 164 to extract users' passwords.
3. In Ignition, create or edit your LDAP directory service and:
 - Select the **Use SSL** check box. The eDirectory server requires an SSL connection in order to use Universal Passwords.
 - Select the **MSCHAPv2 Authentication** check box.

The screenshot shows the 'edir-mv1 - Novell eDirectory Details' configuration window. The 'Settings' section includes the following fields and options:

- Name: edir-mv1
- Service Type: Novell eDirectory
- Security Protocol: Use SSL
- Service Account DN: cn=admin,o=school.edu
- Service Account Password: [masked]
- Directory Root DN: [empty] Browse...
- User Root DN: [empty] Browse...
- Username Attribute: uid Browse...
- Strip Realm
- MSCHAPv2 Authentication
- With LDAP Password Attribute: default Browse...
- With Universal Password

At the bottom, there are sections for 'Primary Server' and 'Secondary Server'.

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:
 - Configure the identity routing policy to use the directory service you saved above.
 - Edit the authentication policy. Make sure it is configured 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.

Procedure

1. In Ignition, create or edit your OID directory service and, in the Directory Server Details window
 - Select the **MSCHAPv2 Authentication** check box.
 - Select the **With Authentication Descriptor** check box.
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.
 - Configure the identity routing policy to use the directory service you saved above.
 - Edit the authentication policy. Make sure it is configured 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.

To test that the OID authentication descriptor attribute is available, 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 231 and [Testing a user lookup](#) on page 190 for instructions.

Creating a Radius Proxy Authentication service

For more information on setting up a Radius proxy server, see [Creating a RADIUS proxy authentication service](#) on page 217.

Managing Directory services

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

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.

Procedure

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.

Name	Directory Type
Internal User Store	Internal Database
Sunnyvale-AD-1	Active Directory
Sunnyvale-LDAP-1	Generic LDAP
MDM01	MDM Service

2. Select the entry for the directory service you want to edit and click **Edit**.

Ignition Server populates the **Directory Service Details** Window with the details of the selected directory service.

3. Edit the details of the selected directory service as required.
4. 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 482 for troubleshooting instructions.

Follow this procedure to rename a Directory Service.

Procedure

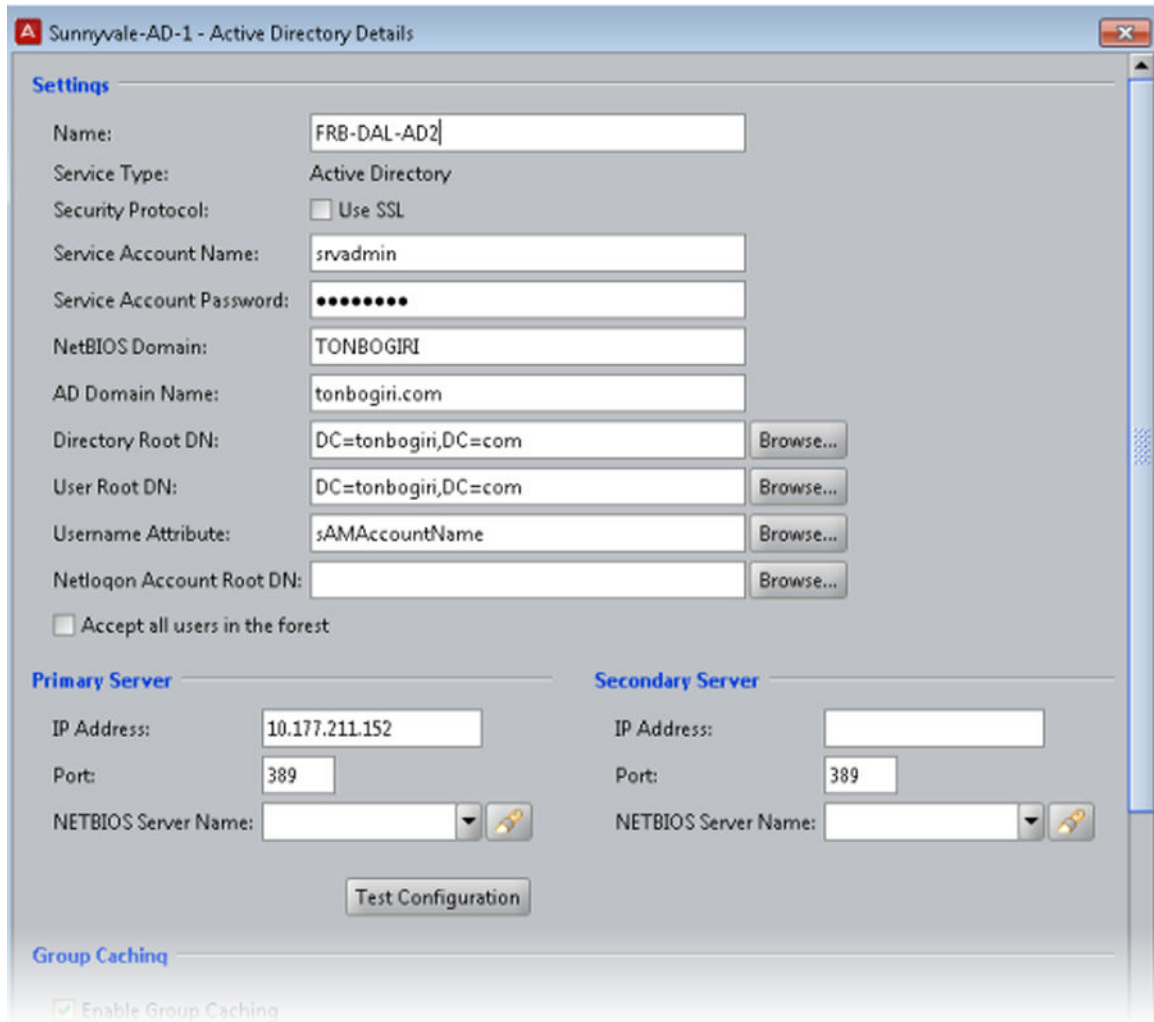
1. In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, and click **Directory Services**.
2. In the **Directory Services** panel, select the entry you want to rename.

3. Click **Edit**.

Ignition Server displays the details for the selected directory service.

4. Enter a different name for the directory service.

5. Click **OK**.



In the preceding sample figure, 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](#) on page 482 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.

Follow this procedure to delete a Directory Service.

Procedure

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 **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 148.

This section explains how to create and manage your directory sets. Before we look at creating a directory set, let's 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.
- 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 described in [Understanding Identity Routing Policy](#) on page 246.

After you create a directory set, you can test it as described in [Checking an Authentication request](#) on page 189.

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 displays.



The directory set 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 and 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

Follow this procedure to add a directory set.

Procedure

1. Click **New** in the **Directory Sets** panel.

The **Add Directory Set** window displays.

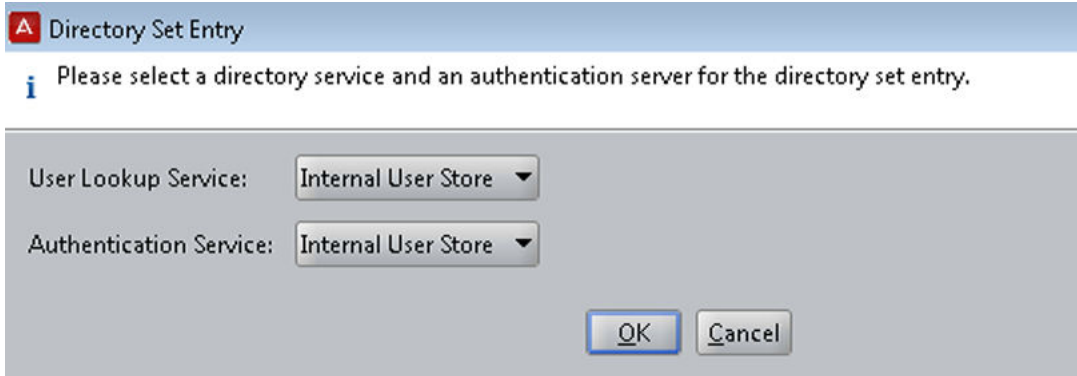


2. Enter an name for the directory set and click **Add**.
3. In the Directory Set Entry window, specify the directory that will provide user account data and group memberships (**User Lookup Service**) and the directory that will authenticate users (**Authentication Service**).


Usually these are one and the same directory. You may choose different directories in cases where you wish to split your authentication from your user lookup, as you might when you couple RSA SecurID authentication with authorization based on AD group membership.

For the example in this document, we will use the internal user store so that we can later demonstrate an authentication of the user account we created earlier. If you have an LDAP or AD user you can test with, then feel free to use your AD or LDAP store, instead:

- In the **User Lookup Service** drop-down list, select **Internal User Store**.
- In the **Authentication Service** drop-down list, select **Internal User Store**.
- Click **OK**.



4. If you are using an AD or LDAP user store, do the following:
 - In the Directory Set window, click **Add** again.
 - In the **User Lookup Service** drop-down list, select the directory service you created earlier. In the example, we use the name `Sunnyvale-AD-1`.
 - In the Authentication Service drop-down list, select your directory service again.
 - Click **OK**.
 - In the directory Set window, click the **Fallthrough** checkboxes in the top row of the table to specify how you want Ignition Server to handle directory failover. By checking these boxes, you can, for example, specify that Ignition Server will attempt authentication against *ActiveDirectory1* if the user's lookup in the *Internal User Store* fails.



5. In the Directory Set window, click **OK** to save the set and dismiss the window.

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.

Procedure

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 248.

Deleting a Directory set

Follow this procedure to delete a directory set.

Procedure

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

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

Adding 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*.

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

Follow this procedure to configure your directory set.

Procedure

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 **Add** in the **Directory Set** panel.
4. In the **Directory Set Entry** window, do the following.
 - a. **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 can 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.

- b. **Authentication Server:** Configure this with the name of the directory service or authentication service that verifies user credentials. In most cases, this is configured with the same name as your User Lookup Service, but many sites that authenticate using RSA SecurID or Kerberos configure this field with the name of the RSA or Kerberos server and configure the **User Lookup** field with the name of an LDAP directory that holds user accounts. See [Authentication service](#) on page 201.

*** 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 211.

- c. 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](#) on page 185. Select the appropriate check box(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**.

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

User Lookup Service	Authentication Service	Fallthrough if Unable to Connect	Fallthrough if User Not Found	Fallthrough if Authentication Failed
TOKUGAWA	TOKUGAWA	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Quicksilver	Quicksilver	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

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

Setting Fallthrough rules to handle lookup and authentication failures

You can configure 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 configured 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.

Authentication Service	Fallthrough if Unable to Connect	Fallthrough if User Not Found	Fallthrough if Authentication Failed
TOKUGAWA	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Quicksilver	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Setting Fallthrough rules

Procedure

1. In the **Directory Set Details** panel, find the row for the directory service whose failure handling you want to configure. The row provides three check boxes, one to configure the fallthrough handling for each of the three cases. The fallthrough rule you configure here applies within the context of this directory set only.
2. In the row, configure the fallthrough check boxes as follows:
 - **Select the check box** 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.
 - **Clear 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 Server

The following table shows how Ignition Server handles fallthrough rules.

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 previous example, 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

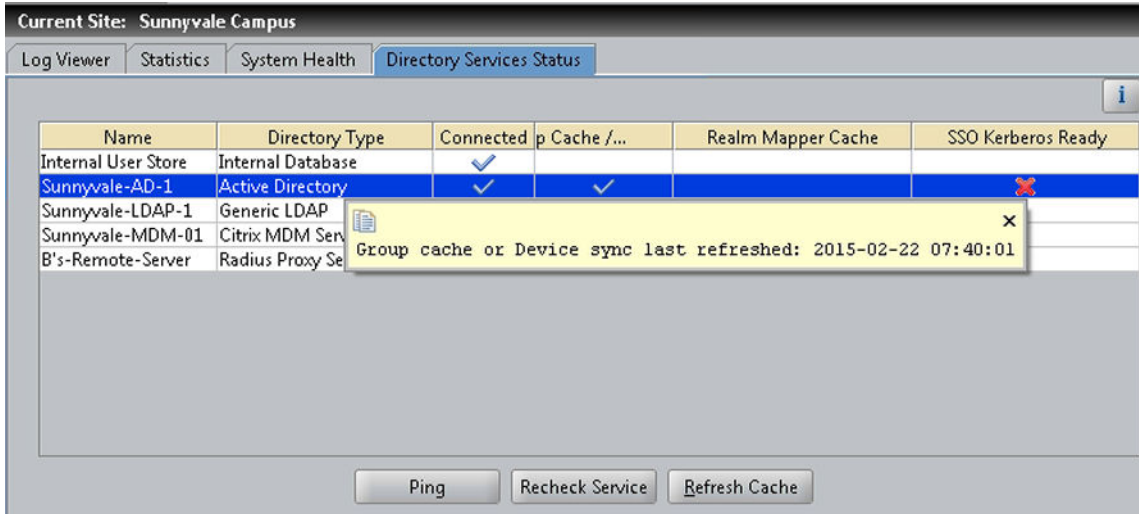
The following sections *do not* cover problems specific to particular types of directory servers. For additional troubleshooting tips related to *Active Directory environments*, see [Problem: Authentication fails on Active Directory](#) on page 484.

Checking directory service connections

Check a directory service connection.

Procedure

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.



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 lets 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
 - 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.

Procedure

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.



Testing a Directory Service connection

Use the **Test Configuration** button to test a directory service connection.

Procedure

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 Configuration** button. Ignition Server tests the connection to the primary directory server and, if configured, the secondary directory server.

! Important:

If Ignition Server returns an error message when you test the connection, see [Problem: Errors occur during Directory Service Set-Up](#) on page 486.

Advanced troubleshooting for Directory Services and Sets

The Advanced Troubleshooting window allows you to test user and device lookups and authentications. The tests that you can perform are.

- [Checking an Authentication request](#) on page 189
- [Testing a user lookup](#) on page 190
- [Testing a Device Lookup](#) on page 190
- [Testing a User Authentication](#) on page 191

! Important:

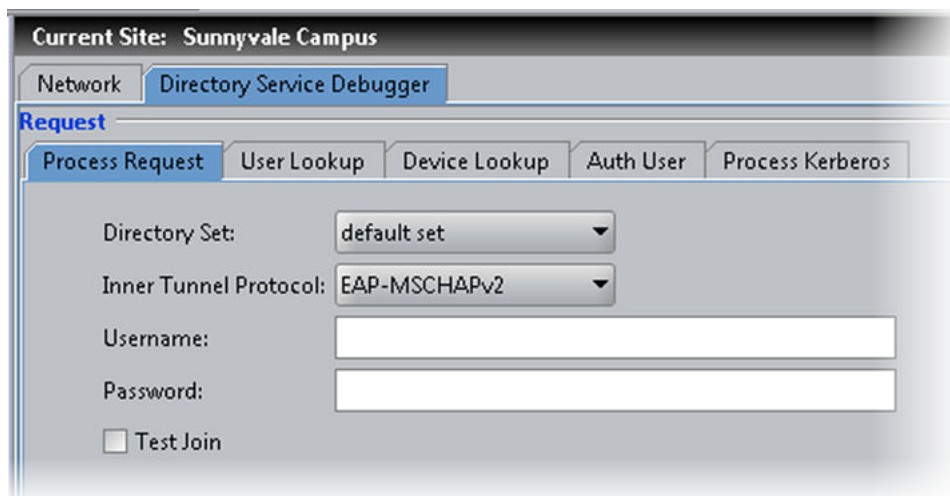
If Ignition Server is running in HA mode, then all test queries originate from the primary node of Ignition.

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:

Procedure

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.



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.

Procedure

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.



5. Type the user's login name in the **Username** field.
6. Click the **Send Request** button to test the authentication and lookup.

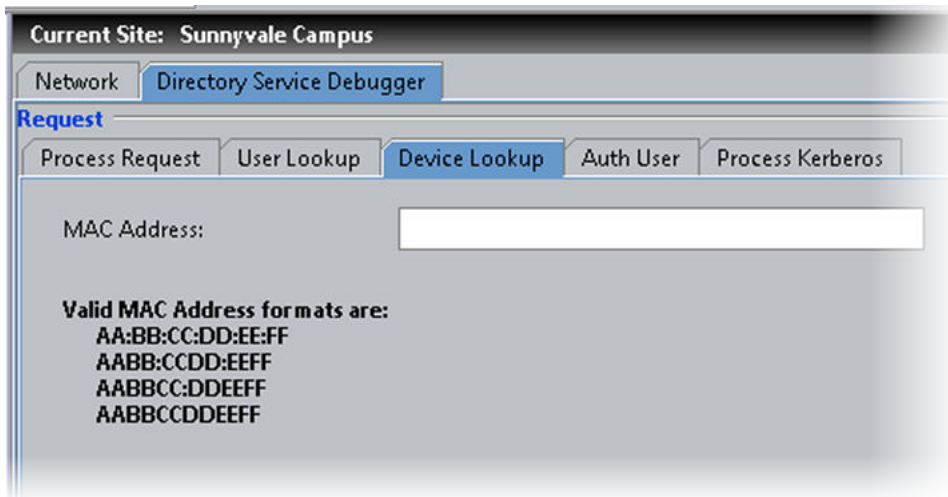
The **Results** field displays the lookup results, and the **Virtual Attributes** displays the attributes retrieved from the user's account record.

Testing a Device Lookup

To find out whether a device is known to Ignition, use the Device Lookup tab of the Advanced Troubleshooting window.

Procedure

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.
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.



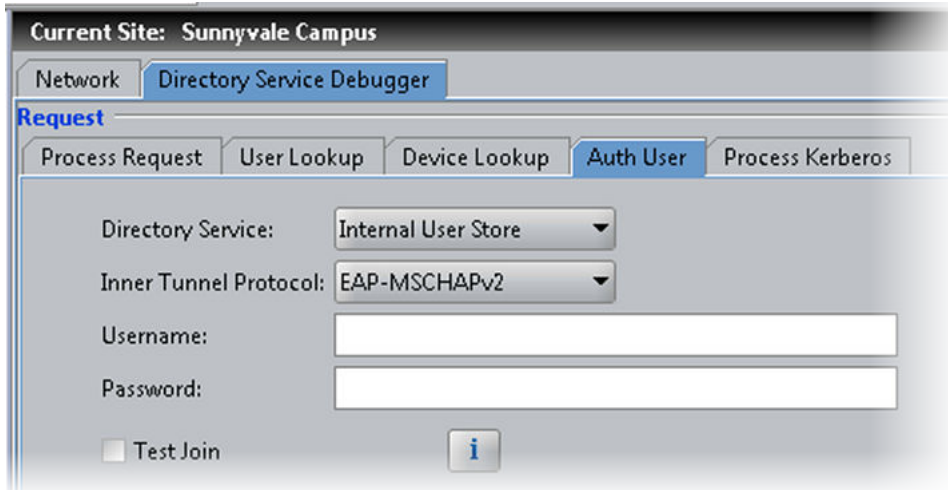
6. Click the **Send Request** button to test the authentication and lookup.

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 133.

Testing a User Authentication

Procedure

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.

Chapter 11: Mobile Device Management

Mobile Device Management (MDM) provides more control for and secure access to bring your own device (BYOD) deployments in a corporate network. With the MDM feature, mobile devices, whether they are corporate-owned or personal, are enrolled in an MDM server. The device attributes such as the OS version are stored in the MDM server and indexed by the MAC address of the device.

The Ignition Server interfaces to the different MDM services to collect the device attributes and save them in the Internal Store. Device attribute lookup happens locally on the Ignition Server. During user authentication, the device attributes are evaluated and fed to the policy engine in the final AAA authorization decision making.

Connecting to an MDM service

You can configure the Avaya Identity Engines Ignition Server to retrieve device attributes from a Mobile Device Management (MDM) server and use them for authorization.

The set of connection settings for the MDM server is called a directory service in Ignition Server. This section shows how to create a directory service for the MDM server.

Creating an MDM directory service

To connect Ignition Server to an MDM server, save the MDM server as a directory service in Ignition Server. The *directory service* specifies the connection settings that Ignition Server uses to connect to the MDM server. You create one directory service for each MDM server you wish to connect to.

Procedure

1. In Dashboard's **Configuration** tab, in the navigation tree, click **Site Configuration**.
2. Click the **Directory Service** link in the main panel.
3. In the Choose Service Type window, select **Citrix MDM Service** and click **Next**.

The Configure Connection to MDM Server window appears.

Configure Connection to MDM Server
Please provide the following information needed to connect to the MDM Server.

Name:

Server Configuration

Server URL:

Server Type:

Username:

Password:

Proxy Server Configuration

Use Proxy Server:

URL:

Username:

Password:

MDM Sync Configuration

Sync on Create:

Resync Duration: (1-168) Hours
Duration after which an auto resync is triggered.


4. On the Configure Connection to MDM Server window, in the **Name** field, enter the name of the MDM service.
5. In the **Server Configuration** section, do the following:
 - a. In the **Server URL** field, enter the URL of the MDM server with instance name.
 - b. In the **Username** field, enter the username of a service account on the MDM server that has Administrator privilege.
 - c. In the **Password** field, enter the password for the service account.

6. If proxy server authentication is required to access the MDM server, in the **Proxy Server Configuration** section, do the following:
 - a. Check the **Use Proxy Server** check box.
 - b. In the **URL** field, enter the URL of proxy server.
 - c. In the **Username** field, enter the username for the proxy authentication.
 - d. In the **Password** field, enter the password for the username.
7. Click **Test Configuration** to make sure that the server information entries are correct.
If a configuration setting is incorrect, Ignition Server warns you. If you receive an error message, correct your settings and test again.
8. In the **MDM Sync Configuration** section, do the following:
 - a. If you want to sync with the MDM server immediately, check the **Sync on Create** check box. To delay the first sync until the Resync interval, do not select this option.
 - b. In the **Resync Duration** field, enter the sync interval, in hours, between the Ignition Server and the MDM server.

This duration does not need to be more frequent than the scheduling time set on the MDM server to synchronize device status.
9. Click **Next**.

The wizard applies your settings to create the directory service in Ignition Server and displays the confirmation page.

Created MDM Server Connection Summary

 Created MDM Server Connection Summary

Name:	MDM01
Server URL:	https://pvmdm01.sv.avaya.com/zdm/ide
Server Type:	CitrixXenMobile
Username:	admin
Password:	***
Proxy Server Configuration	
Use Proxy Server:	No
URL:	
Username:	
Password:	
MDM Sync Configuration	
Sync on Create:	No
Resync Duration:	24

10. If the settings are correct, click **Finish** to create the directory service.
Your directory service has been saved in Ignition Server.

Checking an MDM directory service

To check the MDM directory service, see the following procedures:

- [Checking directory service connections](#) on page 186
- [Checking the Group Cache](#) on page 188
- [Testing a Directory Service connection](#) on page 188

MDM enrolled devices

MDM enrolled devices are device attributes retrieved from Mobile Device Management (MDM) servers and stored locally on the Ignition Server.

MDM Enrolled Devices panel

The **MDM Enrolled Devices** panel lists the devices and attributes learned from the MDM services in the Ignition Server Internal data store. In Dashboard's **Configuration** hierarchy tree, click your site, expand **Site Configuration**, expand **Directories**, expand **Internal Store**, and click **MDM Enrolled Devices** to open this panel to:

- see all MDM enrolled device records
- retrieve a subset of all MDM enrolled devices
- sort and page through MDM enrolled devices
- export MDM enrolled device records

Current Site: Sunnyvale Campus

MDM Enrolled Devices

Get All
 Specify Criteria: MAC Address Starts With

Apply Filter

Viewing records: 1 - 10 of 10

MAC Address	MDM Service	Model	OS Type	OS Version	Registered	Compliant	sk Encrypt...	Pin Lock On	JailBroken
70:6f:6c:69:7...	MDM01	iPad	ios	7.1.2	✗	✓	✓	✗	✓
64:b3:10:b9:...	MDM01	GT-I9082	android	4.2.2	✓	✓	✓	✗	✓
e0:b9:ba:a3:...	MDM01	SYMBIAN	symbian	7.1.2	✓	✓	✓	✗	✗
e0:b9:ba:a3:...	MDM01	Blackberry	bbos	7.1.5	✓	✓	✗	✗	✗
f0:db:f8:14:1...	MDM01	Mobile	windows ph...	7.1.2	✓	✗	✗	✗	✓
f0:db:f8:14:1...	MDM01	PC	PC	8.0	✓	✓	✗	✗	✗
f0:db:f8:14:1...	MDM01	PC	windows rt	8.0	✓	✓	✗	✗	✗
f0:db:f8:14:2...	MDM01	Tablets	windows 8	8.0	✓	✓	✗	✗	✗
01:62:6c:69:7...	MDM01	GT-I9083	windows	8.1	✓	✓	✗	✗	✗
00:23:a9:89:8...	MDM01	PC	windows	8.0	✓	✓	✗	✗	✗

View... Refresh Export...

From this panel you can.

- View the list of all MDM enrolled devices in the internal store.
- Filter the list of MDM enrolled devices to reduce the set of devices to show only those that fit your search criteria.

For information about how to do this, see [Filtering the MDM enrolled devices list](#) on page 197.

- 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 MDM enrolled devices list

Procedure

1. In the **MDM Enrolled Devices** panel window, click **Specify Criteria**.
2. Two drop-down lists are shown to the right of Specify Criteria. In the first list, choose the name of the field you want to filter on. For example, you might choose MAC address, OS Type, OS Version, or JailBroken.
3. In the next drop-down list, select the comparison to be performed. Select **Starts With** or **Equals**.
4. In the text field, enter or select the comparison value.
5. Click **Apply Filter**.

The Dashboard filters the list. To view all devices again, click **Get All**.

Viewing an MDM enrolled device

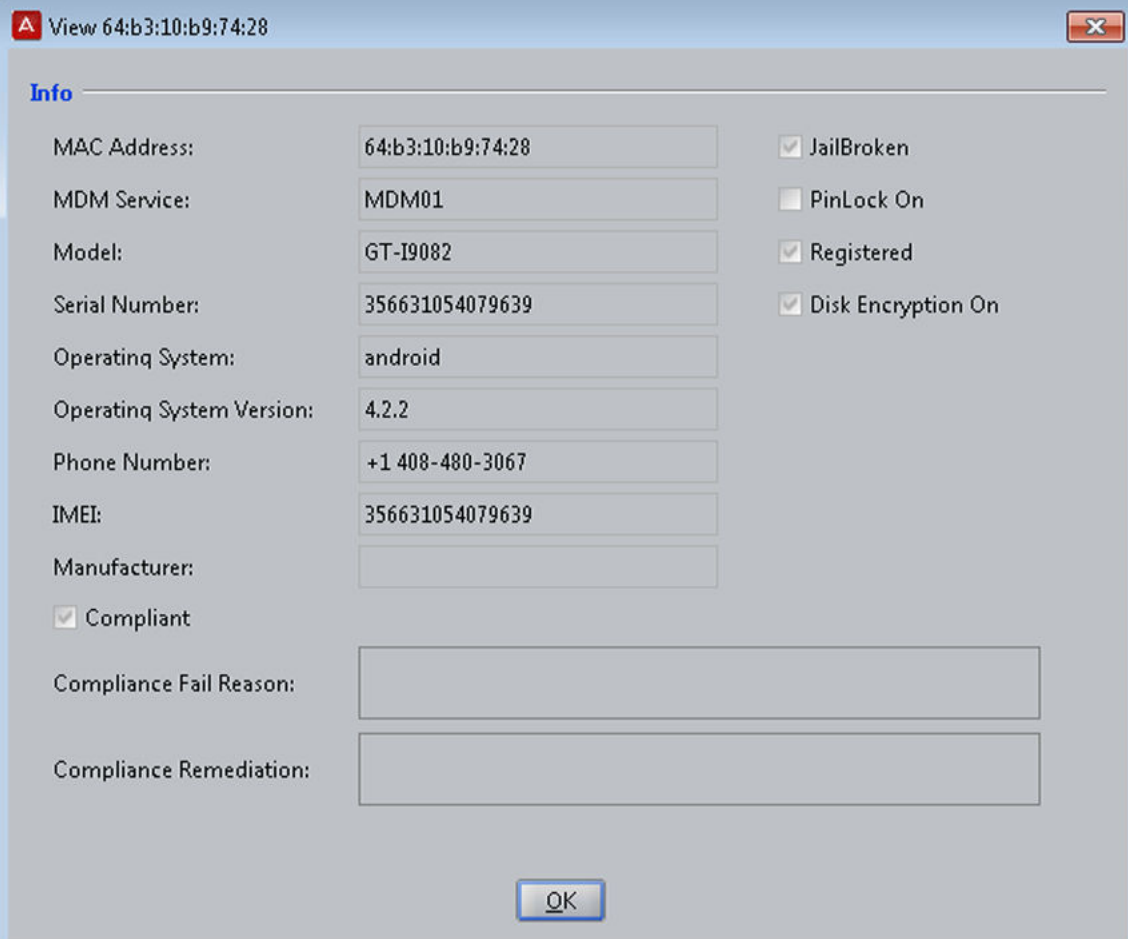
About this task

View the complete details of an MDM enrolled device.

Procedure

1. In the **MDM Enrolled Devices** panel window, click on the desired device entry in the displayed list.
2. Click **View** or double-click on the device entry.

Ignition Dashboard displays the details for the selected device.



The screenshot shows a window titled "View 64:b3:10:b9:74:28" with a close button in the top right corner. The window contains a form with the following fields and controls:

- Info** (Section Header)
- MAC Address: 64:b3:10:b9:74:28
- MDM Service: MDM01
- Model: GT-I9082
- Serial Number: 356631054079639
- Operating System: android
- Operating System Version: 4.2.2
- Phone Number: +1 408-480-3067
- IMEI: 356631054079639
- Manufacturer: (empty field)
- JailBroken
- PinLock On
- Registered
- Disk Encryption On
- Compliant
- Compliance Fail Reason: (empty text area)
- Compliance Remediation: (empty text area)
- OK (button)

3. Click **OK** to close the details for the selected device.

Exporting MDM enrolled device records

Procedure

1. In Dashboard's Configuration hierarchy tree, expand **Directories**, expand **Internal Store**, and click **MDM Enrolled Devices**.
2. In the MDM Enrolled Devices panel, click **Export**.
3. In the MDM Enrolled Device Record Export window, do one of the following.
 - To export all device records, select **Get All**.
 - To export some device records, select **Specify Criteria** and set your filter criteria in the fields 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.

4. Click **Browse** and navigate to find the directory in which you want to save your csv file.
5. Double-click the directory name to select it, type a name for the csv file in the **File Name** field, and click **Save**.
6. In the MDM Enrolled Device Record Export window, click **OK** to export the records.

MDM access policies

You set up MDM access policies to determine which MDM enrolled devices are granted access. The policies are made up of a series of rules that are based on the device attributes.

For more information on using device attributes, see [Using a device attribute in a rule](#) on page 272.

Current Site: Sunnyvale Campus

Access Policy: default-radius-user

Authentication Policy Identity Routing **Authorization Policy**

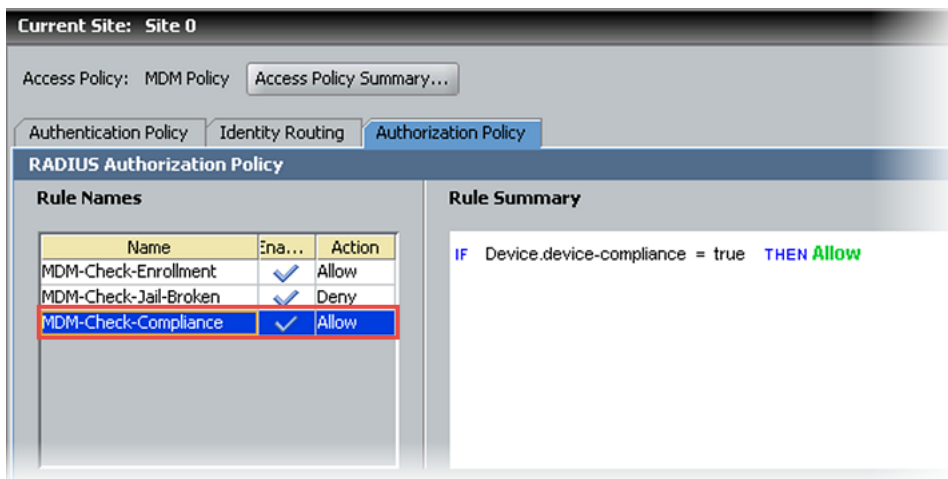
RADIUS Authorization Policy

Rule Names	Rule Summary									
<table border="1"> <thead> <tr> <th>Name</th> <th>Enabled</th> <th>Action</th> </tr> </thead> <tbody> <tr> <td>Allow-All</td> <td>✓</td> <td>Allow</td> </tr> <tr> <td>mdrn-policy</td> <td>✓</td> <td>Allow</td> </tr> </tbody> </table>	Name	Enabled	Action	Allow-All	✓	Allow	mdrn-policy	✓	Allow	<pre>IF { Device.device-registered = true AND Device.device-disk-encryption = true AND Device.device-jailbroken = false } THEN Allow</pre>
Name	Enabled	Action								
Allow-All	✓	Allow								
mdrn-policy	✓	Allow								

The following example shows an access policy that only checks if a particular device is enrolled in the MDM. Device.source is the name of the MDM service.



The following example shows an access policy that only checks if a particular device is in compliance with the policy on the MDM server.



Chapter 12: Authentication service

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.

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.

Procedure

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 486.
 - 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.

Overview of Token Authentication in Ignition Server

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.

How a user logs in with Token Authentication

Ignition Server handles a token authentication login as follows.

Procedure

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 *nexttokencode*. 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 Server

This section explains how to configure Ignition Server to require token authentication.

Prepare your Token Server

Install and configure your token server software and its required supporting software. Note the following tips.

1. Configure DNS settings on your token server so that it can resolve the address of the Ignition Server RADIUS service.
2. Configure your token server to recognize Ignition: In the RSA Authentication Manager configuration, add an *Agent Host* record to represent the Ignition Server. For other types of token servers, 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 configure 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 Servers, 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.

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.

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.

Procedure

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 68.
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. 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 220.

! **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 using 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 204.

To set up a PAP RADIUS connection with a token server, define your token server in Ignition Server as an *authentication server*, as described here.

Procedure

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 68.
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 486 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.

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.

Procedure

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 148.
4. In the **Authentication Service** drop-down list, select the name of your RSA SecurID authentication service or token authentication service.

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.

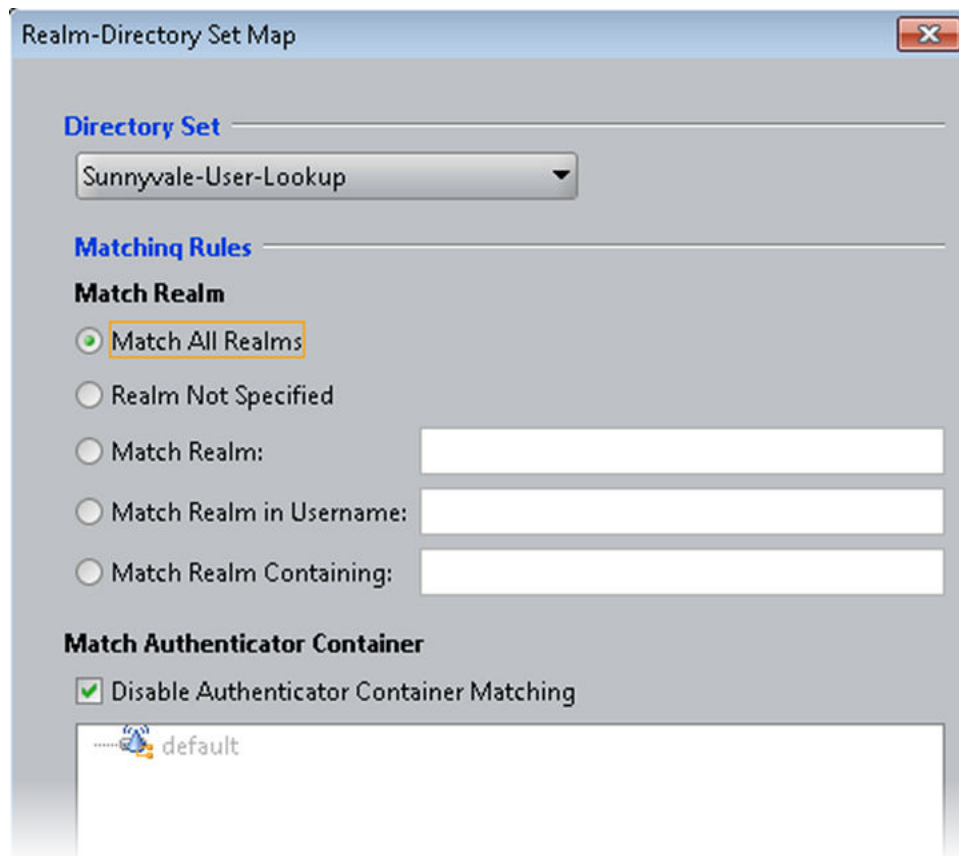
Set Up an Access Policy that uses Token Authentication

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 you are connecting to the token service through PAP), and it must use the Ignition Server directory service that includes your token server. Configure this as follows.

Procedure

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 select 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 through 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.

Prepare Your Authenticators

Your authenticators are your switches or VPN concentrators that send authentication requests to Ignition. In each such switch or VPN concentrator, configure Ignition Server to act as the RADIUS Server. Do the following.

Procedure

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 Servers, 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 through RADIUS.
3. **DNS server:** Configure DNS settings on each authenticator so that it can resolve the address of the Ignition Server RADIUS service

Example

- **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-3DESMD5*, 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.

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. See [Creating an authenticator](#) on page 104.

Make sure token-capable clients are installed

Each end user who authenticates 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.

Perform an End-to-End test

Send an authentication request to test your configuration.

Procedure

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.

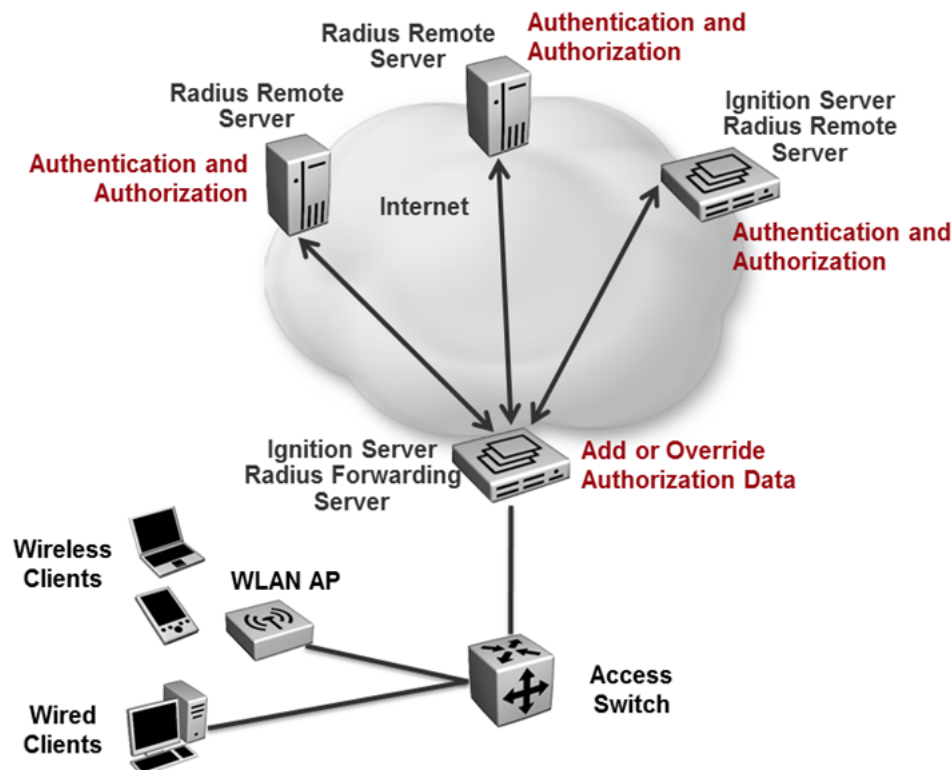
Troubleshooting Notes: If the RSA service initializes correctly but all authentications fail, and the RSA server reports that the node secret has been sent to the client (the message “node validation failed” appears in RSA server logs), there is a possibility that the node secret is out of sync. Follow the steps in [Handling changes to the Node Secret](#) on page 220 to re-create the service.

Setting up a RADIUS proxy server

A RADIUS proxy server forwards RADIUS requests to a remote 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 forwarding server performs local authorization after receiving a response from the remote server to suit the local network deployment. After the forwarding server completes authentication, the information is logged for both success and failure.

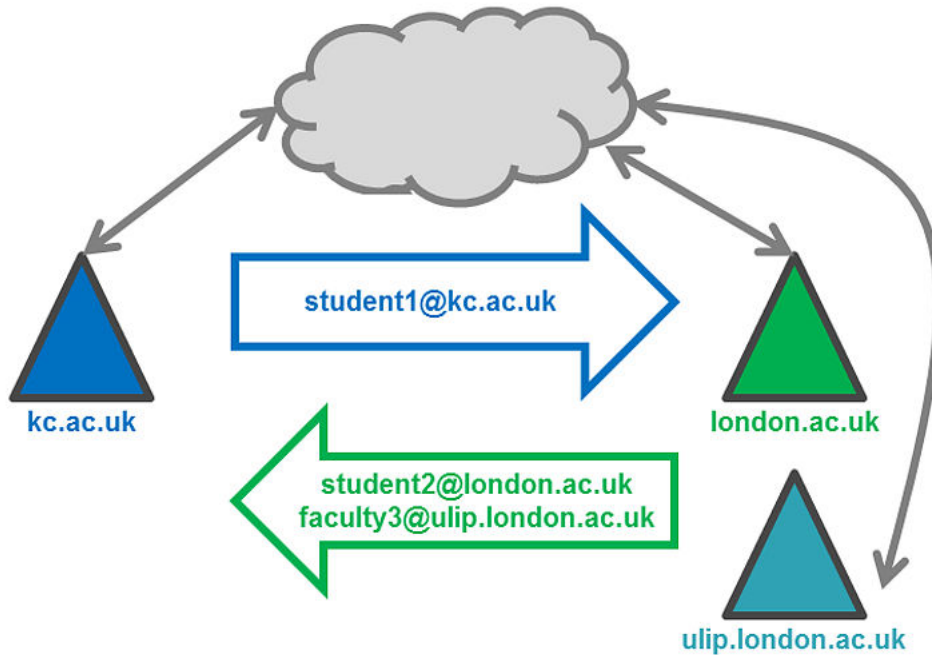
If you are using a RADIUS proxy server, you must configure an authentication service in Ignition Server. In Ignition Server, you manage authentication services in the Directory Services panel, in the same way you manage directory services.




Use case examples

The following diagrams show several different examples of Advanced RADIUS Proxy deployments.


In this example, Student1 belongs to Kings College (kc.ac.uk) and visits London College (london.ac.uk). Student2 belongs to London College (london.ac.uk) and visits Kings College (kc.ac.uk). Faculty3 belongs to University of London in Paris (ulip.london.ac.uk domain) and visits Kings College (kc.ac.uk).



In this example, Engineer1 belongs to us.avaya.com. Engineer1 visits the EMEA office and tries to connect with the local network emea.avaya.com. Engineer1 can be authenticated against the local LDAP in us.avaya.com and authorized at emea.avaya.com based on local enforcements.



In this example, usaguest is a contractor working for Avaya US and is given access to us.avaya.com. usaguest travels to the Avaya office in Germany, and tries to connect to the domain emea.avaya.com. usaguest is authenticated remotely at us.avaya.com and authorization policies are set based on what is defined in emea.avaya.com.



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 181.

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**. 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 directories and authentication servers to a directory set](#) on page 183.

Creating a RADIUS Access Policy for RADIUS Proxy Server

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](#) on page 237.

Creating a new RADIUS Proxy Policy

Use this procedure to create a new RADIUS Proxy Policy and add authorization policy rules.

Each rule consists of one or more constraints. Each constraint tests the value of an attribute. If there are multiple constraints, you can join them into separate logical statements to ensure the proper order of authorization as required.

The rule action determines whether the user is denied or granted access based on the defined constraints.

Procedure

1. In Dashboard's **Configuration** hierarchy tree, expand **Access Policies** and click **PROXY**. Click **New**.
2. Enter the **Access Policy Name** and click **OK**.
3. Highlight the new access policy name, and click **Edit**.

The Edit Authorization Policy window displays.

4. Do one of the following:
 - To add a new rule, click **Add** in the Rules panel, enter a **Name** for the new rule and click **OK**.
 - To copy an existing rule, click **Copy** in the Rules panel, select the desired rule, and click **OK**.
5. To set up rule details, highlight the rule name in the **Rules** list.

The rule details are shown in the **Selected Rule Details** pane. Any existing constraints for the selected rule are listed in the **Constraints** list.

6. Do one of the following:
 - To add new constraints, click **New**.
 - To edit existing constraints, highlight the constraint and click **Edit**.
7. From the **Attribute Category** drop-down list, select the category.
All of the valid attributes for the category are listed.
8. Select the desired attribute.
The configurable details for the selected attribute are displayed.
9. Configure the attribute details as applicable:
 - Select the comparison operator.
 - Select the format.
 - To compare the attribute value with a fixed value, select the **Static Value** radio button and type or choose the comparison value in the field below.
 - To compare the attribute value with a value retrieved from another attribute, select the **Dynamic Value of Attribute** radio button. In the drop-down list below, choose the Attribute Category. In the second drop-down list, choose the attribute that should provide the comparison value. The list of comparison attributes contains only those attributes whose data type matches the data type of the constraint attribute.
10. Click **OK**.
11. Repeat Steps 6 through 10 for each constraint.
12. To logically group multiple constraints, in the **Constraint** list, highlight the first and last constraints to be grouped and use the opening and closing parentheses drop-down lists to group the constraints. Use the **AND/OR** drop-down list to form a logical condition statement.
13. Do one of the following:
 - Select **Deny** for the **Action** and go to Step 15.
 - Select **Allow** for the **Action**.
14. If you chose **Allow** for the **Action**, do the following:
 - In the **Send Attributes** row, click the Edit icon, and use the left and right arrows to add or delete attribute values from the **Attribute List**.
The forwarding server updates (if present) or adds (if not present) these attributes to the remote server response before sending to the authenticator.
 - In the **Delete Attributes** row, click the Edit icon, and use the left and right arrows to add or delete attribute values from the **Attribute List**.
The forwarding server deletes these attributes from the remote server response before sending to the authenticator.

Note that, when a forwarding server receives a response from a remote server, the first Delete Attribute is applied, and then the second, and so on. All of the attributes defined in the Delete Attribute List on the forwarding server are deleted first. After that, the first Send

Attribute will either add the attribute or update an existing attribute value that may be present in the remote server response. Then the second, and so on. After applying Delete, Send (in that order), the forwarding server sends a response back.

15. Check the **Summary** section to confirm the rule details, and click **OK**.

The policy and associated rules is saved.

Creating a RADIUS proxy authentication service

Use this procedure to create a RADIUS proxy authentication service. The Create Service Wizard guides you through the steps.

Procedure

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.
4. Enter the **Shared Secret** for the RADIUS proxy server.
5. Select the **Proxy Policy** from the drop-down list.

This policy determines how to update the RADIUS response from the remote server and change the authorization attributes to suit the local network deployment. This policy can only be associated with the Radius Proxy type of directory services and include only authorization.

The list contains the proxy policies configured on the system. By default, it is associated with a default policy that has no local authorization.

For more information about configuring the proxy policies, see [Creating a new RADIUS Proxy Policy](#) on page 215.

6. 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.

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 reachability.

With Keepalive turned on, Ignition Server periodically looks up the supplied username/ password on the remote server to determine reachability, and if successful, marks the service as *Connected* in the **Directory Services Status** tab. By default, Ignition Server uses a predefined username and 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.

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 the 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 that the remote server is always connected and the requests may get dropped if the remote server health status is not determined.

7. Specify the **IP Address** and **Port** for the primary RADIUS proxy server and optionally for the secondary RADIUS proxy server.

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.

8. Click the **Test Keepalive** button.

Testing the connection may take a few minutes. If a configuration setting is incorrect, Ignition Server warns you.

9. Click **Next**.

The next window summarizes the connection settings of the service.

10. Click **Finish**.

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

Configuring the remote RADIUS server

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

To edit your connection to an authentication service, use the following procedure.

Procedure

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, edit the details of the service as required.
 4. 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.

Procedure

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 220 to re-create it.

Procedure

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 do the following.

Procedure

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 204.
- 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.

Set up the supplicant and authenticator for SecurID authentication.

Procedure

1. Set your supplicant to use EAP-GTC authentication. In OpenSEA's XSupplicant, create a **Profile** that uses the **Tunnel Protocol EAP-GTC**.
2. 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.
3. 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.

Chapter 13: 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.

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
- Mobile Device Management (MDM) 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 [User authentication policy](#) on page 237.

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.

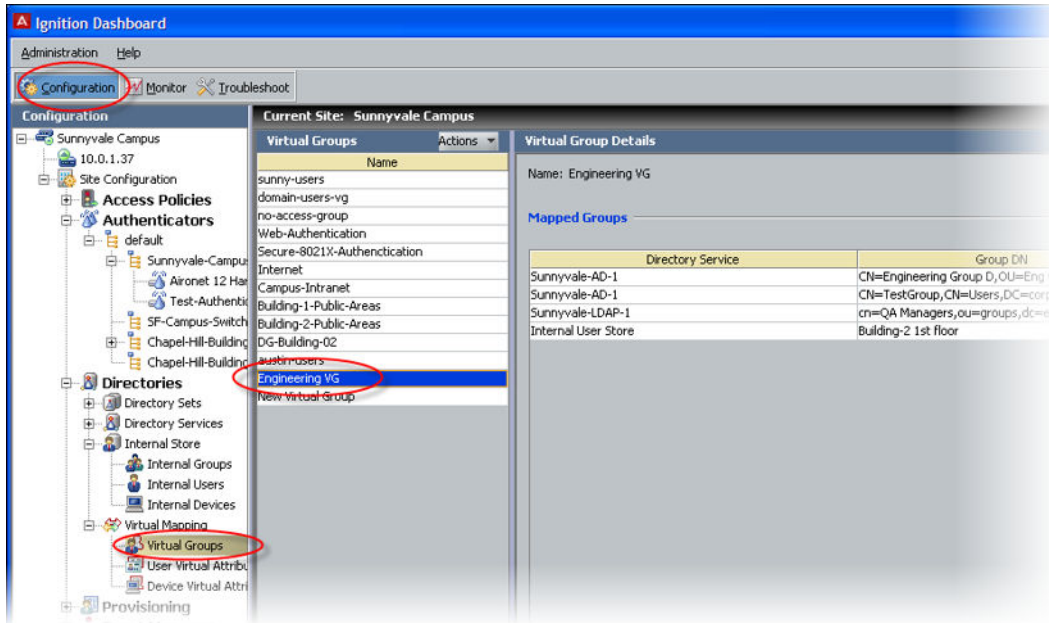
Important:

Ignition Server provides a similar mechanism for abstracting attribute names. See [User Virtual Attributes](#) on page 229.

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.



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 following table 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.

	AD	Sun	Open-LDAP	Novell eDirectory	Oracle OID
Group saved as an organization or organizationUnit entry	Yes	Yes	Yes	Yes	Yes
Group saved as a groupOfNames or groupOfUniqueNames entry with members listed in member	No	Yes	Yes	Yes	Yes

Table continues...

	AD	Sun	Open-LDAP	Novell eDirectory	Oracle OID
or uniqueMember					
Group listed in the user's record, in the nsRole attribute (object class is ldapsubentry)	No	Yes	No	No	No
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

Checking 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

Follow this procedure to add a new virtual group.

Procedure

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

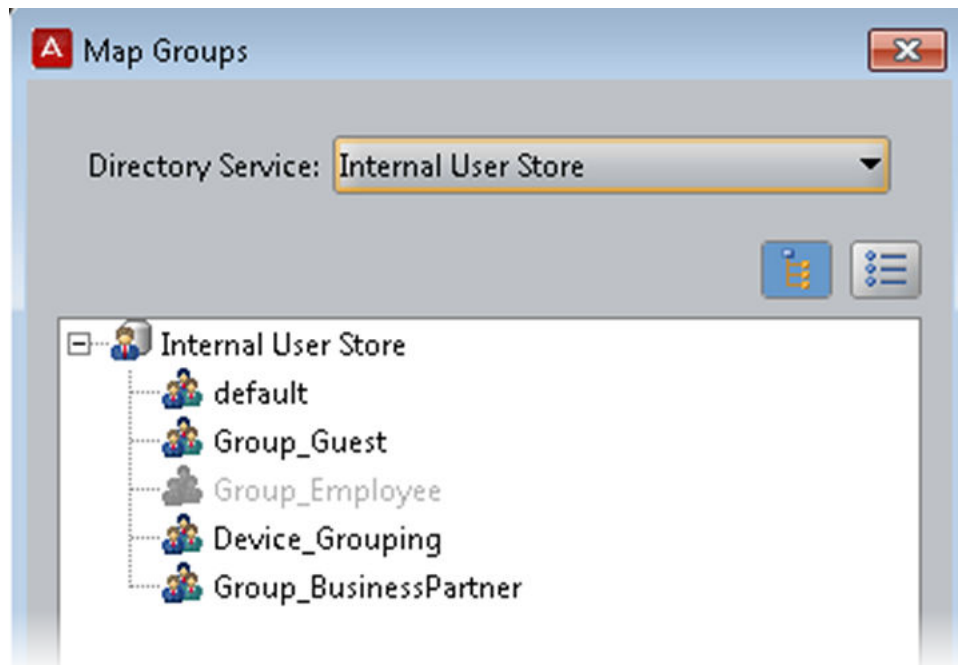
You can map groups from your directory services onto this new virtual group.

Follow this procedure to map groups from available directory services

Procedure

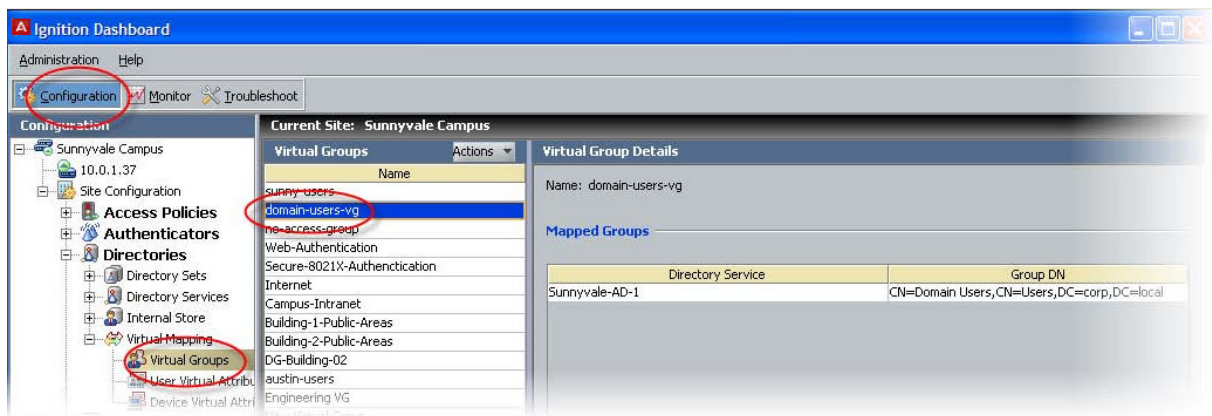
1. Select the virtual group from the displayed list of virtual groups in the **Virtual Groups tab**.
2. Click **Add** in the **Virtual Group Details** section of this window.

The **Map Groups** window displays.



3. Select the **Directory Service** from the drop-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 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**.



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 482.

Procedure

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 482.

Procedure

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. 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 283.

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.

! Important:

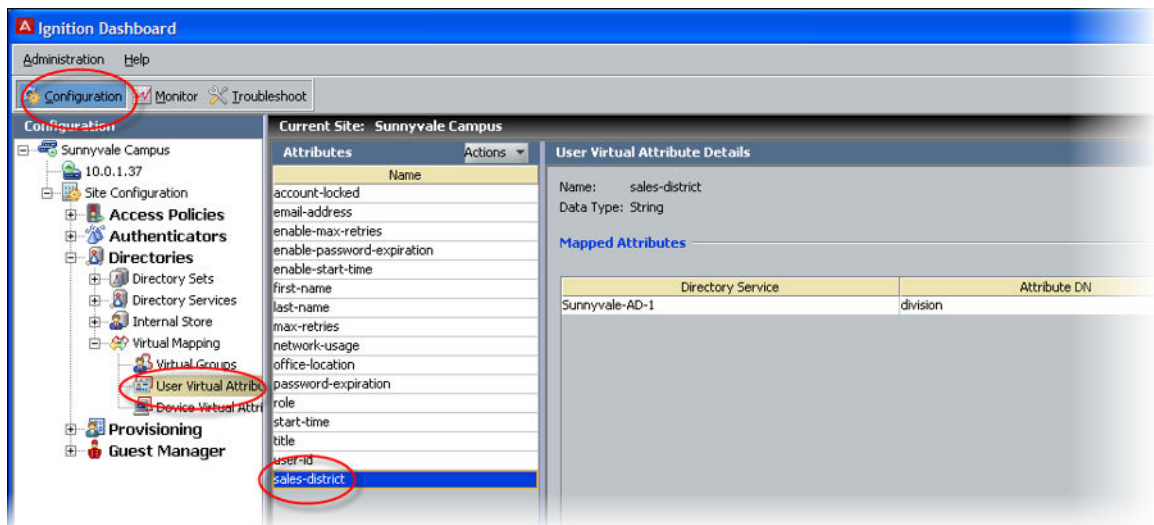
Ignition Server provides a similar mechanism for abstracting group names. See [Virtual Groups](#) on page 224.

Browsing User Virtual Attributes

To view 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 displays.



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 can 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

Add a new user virtual attribute.

Procedure

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 351.
 - **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 [Managing Directory services](#) on page 176.

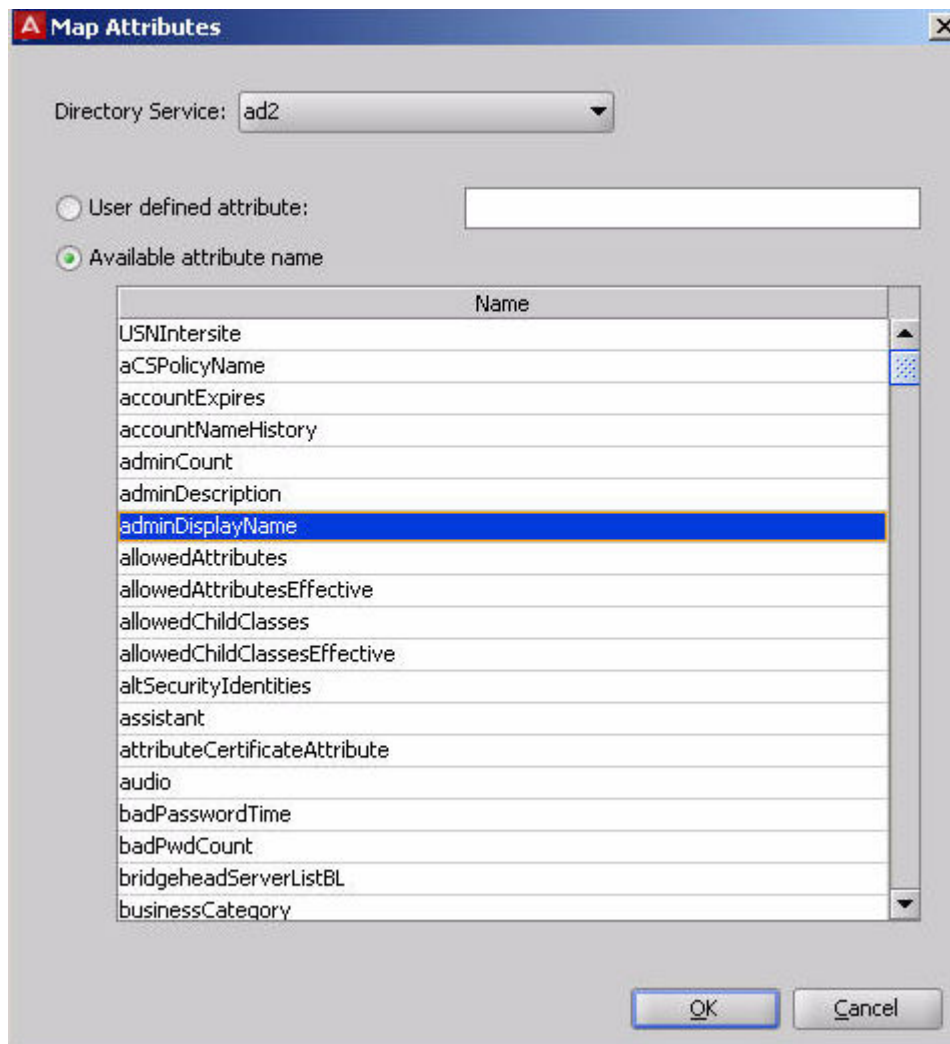
Mapping Directory Service Attributes to User Virtual Attributes

Map attributes (Distinguished Names) from a directory service object to a user virtual attribute.

Procedure

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.



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 188.

! **Important:**

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 190.

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 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 283.

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 482 for troubleshooting instructions.)

Procedure

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**.

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 482 for troubleshooting instructions.)

Procedure

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.

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.

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.

Procedure

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).

The screenshot shows a web interface for editing a device record. The title bar reads "b9:ff:8b:5a:7a:3e - Device Record Details". Under the "Info" tab, the following fields are visible:

- MAC Address: b9:ff:8b:5a:7a:3e
- Name: HP-Laserjet-Floor2
- Type: printer
- Source: devicestoimport.csv
- VLAN Label: hq-printer-vlan
- VLAN ID: 206

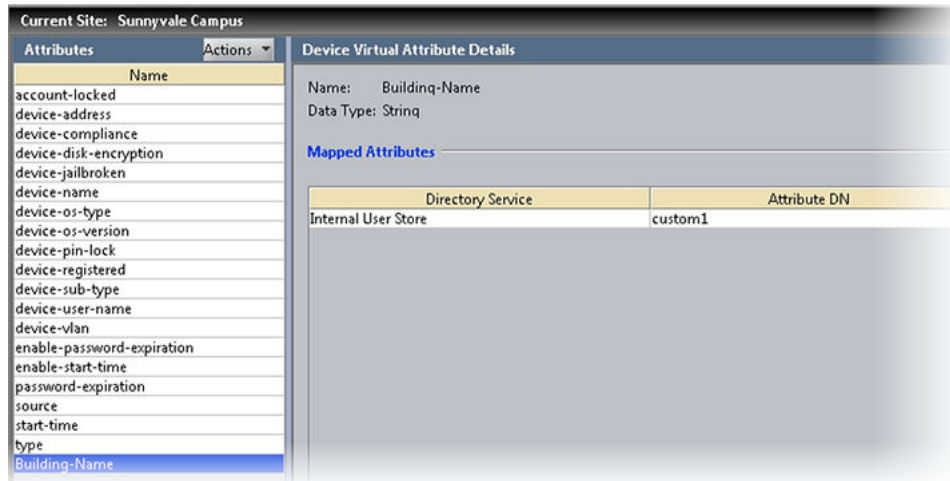
Below the "Info" section is the "Custom Attributes" section, which contains six input fields labeled "custom 1" through "custom 6". The "custom 1" field contains the text "Sunnyvale-Building-2" and is highlighted with a red circle.

2. Create your virtual attribute:
 - In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, expand Directories, expand Virtual Mapping, and click **Device 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*.



For instructions on using the attribute in an authorization rule, see [Using a device attribute in a rule](#) on page 272.

Chapter 14: 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.

Additional policy topics

Subsequent sections of this document cover the remaining topics in access policy.

- [User Authorization Policy](#) on page 252 explains how to set policies that control user access based on user attributes, transaction details, and network hardware specifications.
- [Provisioning policy](#) on page 274 explains how to set policies that assign users to VLANs and/or set authenticator attributes.
- [VLAN assignment using the Device Record VLAN fields](#) on page 348 shows how to set up provisioning policies that assign each user to an appropriate VLAN.
- [Authentication service](#) on page 201 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 316 explains how to write a policy that requires each connecting device to have a valid Active Directory account.
- [Introduction to MAC Authentication](#) on page 339 shows how to set up MAC address checking and how to set up an access policy for devices incapable of 802.1X authentication.
- [Introduction to Asset Correlation](#) on page 353 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 97.

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 246.
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 240.
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 253.
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 301.

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."

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.

The screenshot shows the Ignition Dashboard interface. The left sidebar displays the Configuration hierarchy tree, with 'Sunnyvale-RADIUS' selected under 'Access Policies'. The main panel shows the configuration for 'Sunnyvale-RADIUS-policy' under 'Current Site: Sunnyvale Campus'. The 'Authentication Policy' tab is active, showing a table of rules for 'Authenticated RADIUS Authorization Policy'.

Callouts provide the following information:

- Selecting an access policy displays the settings for its authentication, identity routing, authorization, and provisioning policies**: Points to the left sidebar where 'Sunnyvale-RADIUS' is selected.
- Allows you to specify which authentication protocols are permissible**: Points to the 'Authentication Policy' tab.
- Allows you to specify how the user account is looked up**: Points to the 'Identity Routing' tab.
- Click to edit the currently selected tab**: Points to the 'Edit...' button next to the 'Authenticated RADIUS Authorization Policy' tab.
- User authentication rules**: Points to the table of rules.

Rule Names	Name	Ena...	Action
	Check...	<input checked="" type="checkbox"/>	Allow
	Allow...	<input checked="" type="checkbox"/>	Allow

Rule Summary: IF User group-member is any one of [domain-users-vg] THEN Allow

If No Rules Apply: Deny

Unauthenticated RADIUS Authorization Policy - Currently Disabled

The authorization policy is typically not executed unless authentication succeeds. Define and enable this policy if you want to authorize users even when they fail to authenticate.

Enable Unauthenticated RADIUS Authorization

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 240.

- **identity routing policy** specifies where Ignition Server can find user records and how it should handle user lookup failures. Identity routing policies are described in [Understanding Identity Routing Policy](#) on page 246.

If authentication succeeds, then Ignition Server executes the user *authorization* policy. See [User Authorization Policy](#) on page 252 .

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.

The follow tables lists Non-EAP Authentication Protocol and User Data Store compatibility.

Data Store Type	NONE/PAP	NONE/CHAP	NONE/MSCHAPv2	TTLS/PAP
Ignition Server internal	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
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 170.

**The Ignition server acts as a proxy and forwards requests to the remote proxy server for authentication.

The follow tables lists EAP Authentication Protocol and User Data Store compatibility.

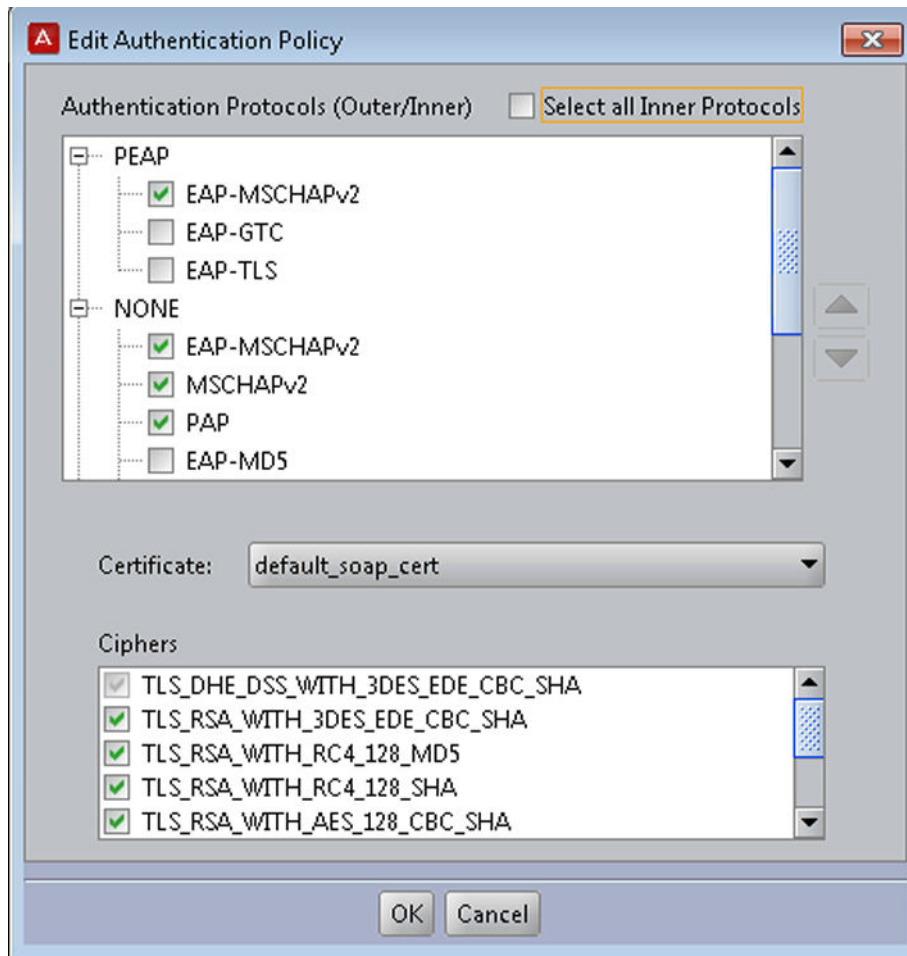
Data Store Type	NONE/EAP-MD5	NONE/EAP-GTC	NONE/EAP-TLS	NONE/EAP-MSCHAPv2	PEAP/EAP-GTC	PEAP/EAP-MSCHAPv2	PEAP/EAP-TLS
Ignition Server internal	Yes	*	Yes	Yes	*	Yes	Yes
Microsoft Active Directory		*	Yes	Yes	*	Yes	Yes
Sun Java System Directory Server (SunONE LDAP)		*		Yes**	*	Yes**	

Table continues...

Data Store Type	NONE/ EAP-MD5	NONE/ EAP-GTC	NONE/ EAP-TLS	NONE/ EAP- MSCHAPv 2	PEAP/ EAP-GTC	PEAP/ EAP- MSCHAPv 2	PEAP/ EAP-TLS
Novell eDirectory		*		Yes**	*	Yes**	
Oracle OID		*		Yes**	*	Yes**	
Generic LDAP		*		Yes**	*	Yes**	
RSA Authentication Server		Yes			Yes		
<p>* 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.</p> <p>** 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 170.</p>							

Authentication Policy window

Use the Authentication Policy window to establish an allowed set of user authentication protocols for the access policy.



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.

- user authentication
- device authentication using Windows machine authentication. (See [Windows Machine authentication](#) on page 316.)

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 83), and can be shared across access policies.
- A potential defect in Microsoft Windows XP prevents the use of DSAsigned 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

Follow this procedure to configure the authentication policy of your access policy.

Procedure

1. In Dashboard's Configuration hierarchy tree, expand Access Policies and expand RADIUS.

- Click the name of your access policy and click the **Authentication Policy** tab. Click **Edit**.
- 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.
- Select the check box next to each desired authentication type. Choose as many as you want to support. See [Supported Authentication Types](#) on page 240 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 92.

- 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 preceding steps for each outer tunnel type.
- 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 91.

! **Important:**

For tips on choosing a certificate, see [Factors that limit your choice of a Protocol Credential Certificate](#) on page 244.

- 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.)

- Click **OK**.

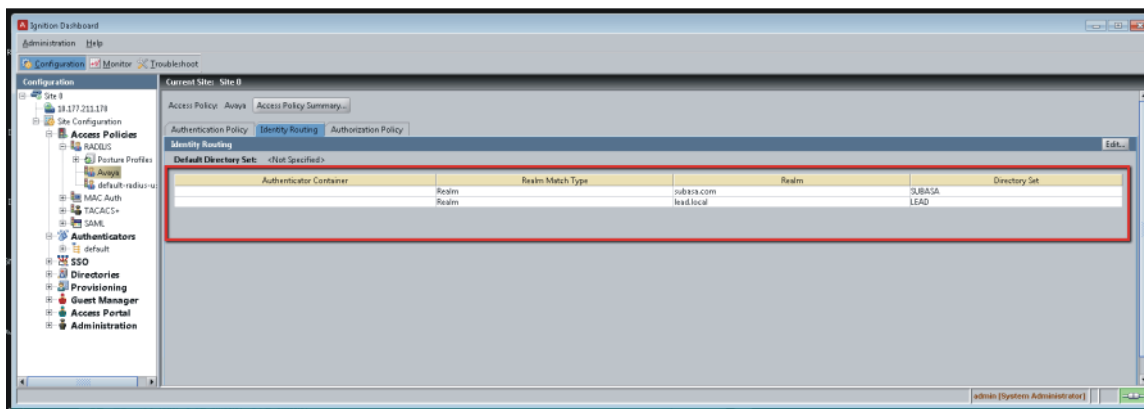
Your authentication policy has been configured.

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. see [Understanding Identity Routing Policy](#) on page 246.

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.



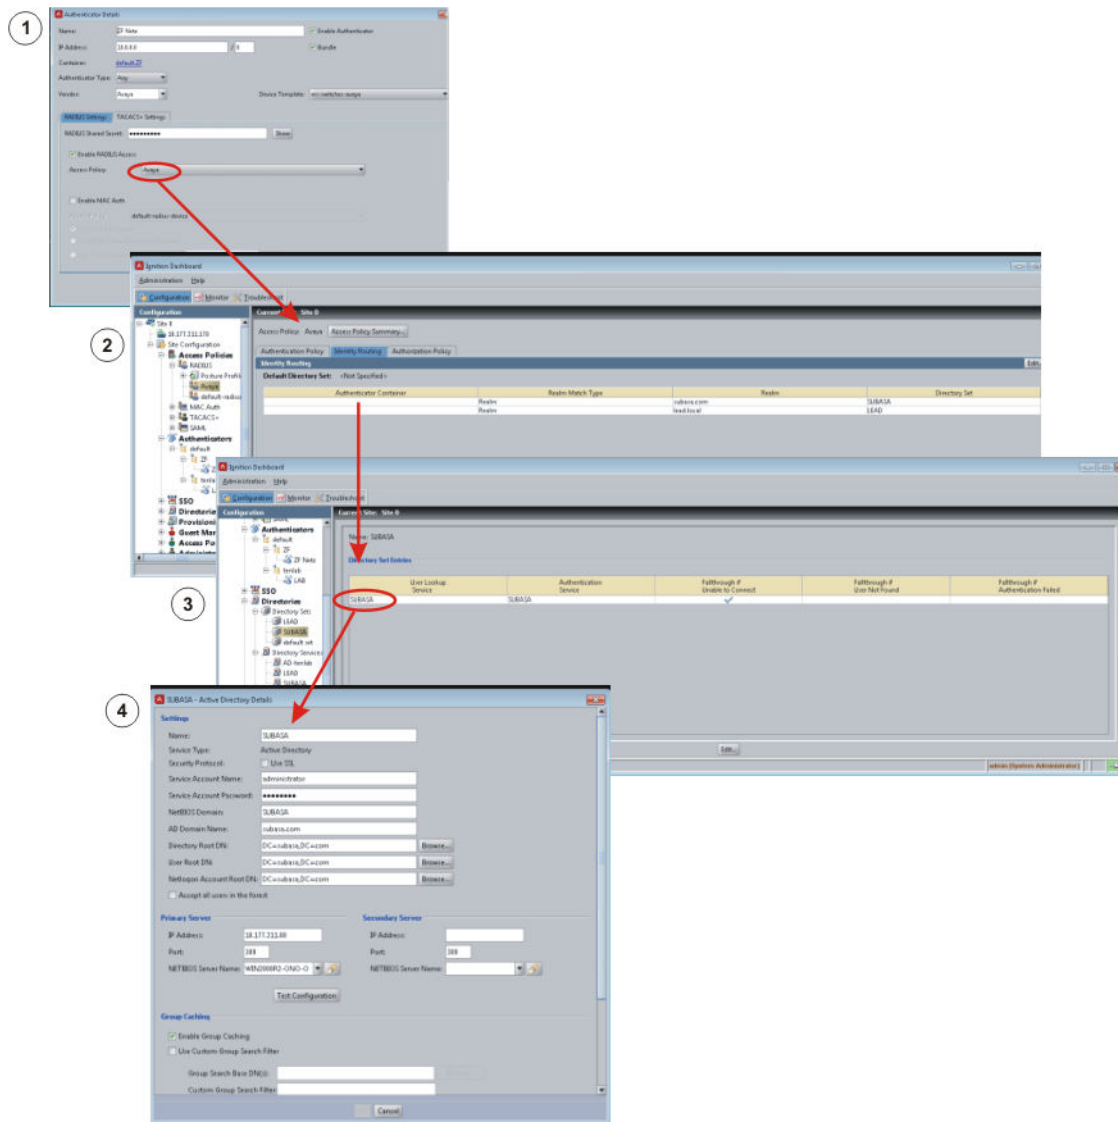
How Ignition Server looks up a user for Authentication and Authorization

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 97 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 189.)

The following example shows 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. Use the following procedure to create your identity routing policy.

Procedure

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 **Edit**.
3. If you want to enable a default directory set, select the **Enable Default Directory Set** check box and select the name of your default directory set in the drop-down list just below the check box.

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. Configure your realm/authenticator mapping rules in the **Realm-Directory Set Mapping** table. To begin adding a mapping rule, click **New** 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.

The screenshot shows the 'Realm-Directory Set Map' dialog box. It has a title bar with a close button. The 'Directory Set' section has a dropdown menu currently showing 'default set'. The 'Matching Rules' section is divided into two parts. The 'Match Realm' section has four radio button options: 'Match All Realms' (which is selected and highlighted with a yellow box), 'Realm Not Specified', 'Match Realm:' (with an empty text field), 'Match Realm in Username:' (with an empty text field), and 'Match Realm Containing:' (with an empty text field). The 'Match Authenticator Container' section has a checked checkbox labeled 'Disable Authenticator Container Matching'. Below this is a tree view showing a folder icon next to 'default', which contains two sub-items: 'Chapel-Hill-Building-1' and 'Chapel-Hill-Building-2', each with a folder icon. At the bottom of the dialog are 'OK' and 'Cancel' buttons.

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, configure your realm-matching rule.
 - To create a rule that requires a realm-name match, select the **Match Realm** check box 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 such as *jlee@avaya.com*, then you specify *avaya.com* here. If they log in with names such as *avaya/jlee*, specify *avaya* here.

*** Note:**

The comparison is case-sensitive.

For more information see [Additional notes on Realm-Matching rules](#) on page 251.

- To create a rule that matches realm-less user names (for example, a user name of *jlee*), select the **Realm Not Specified** check box.
- To specify that no realm-matching is required, select the **Match All Realms** check box.

The rule must match all user names. This is useful if you want to perform authenticator-matching but no realm-matching.

- To create a rule that matches a realm in a username, select the **Match Realm in Username** check box and type the username in the neighboring field.
- To create a rule that does a partial match of a realm, select **Match Realm Containing** and enter the realm in the field.


For example, if you enter *avaya.com*, requests coming from *jsmith@avaya.com*, *jsmith@ca.avaya.com*, or *jsmith@us.avaya.com* are all found. This option is especially useful when the users in an organization are distributed into multiple realms/sub-domains, as you do not have to add rules for each sub-domain or realm.

7. In the **Match Authenticator Container** section, configure 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, select the **Disable Authenticator Container Matching** check box.
- To use authenticator-matching, make sure the check box is *not* selected, 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 98.

8. Click **OK**.
9. In the **Identity Routing Policy** window, add additional realm/ authenticator mapping rules by clicking **New** again and repeating Step 6.



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 at the right of the Realm Directory Set Mapping table.
11. Click **OK** to dismiss the Identity Routing Policy window.

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*.

Chapter 15: 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 274.

Ignition Server also lets you define authorization policies for devices. See [Introduction to MAC Authentication](#) on page 339.

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.

Important:

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 271.

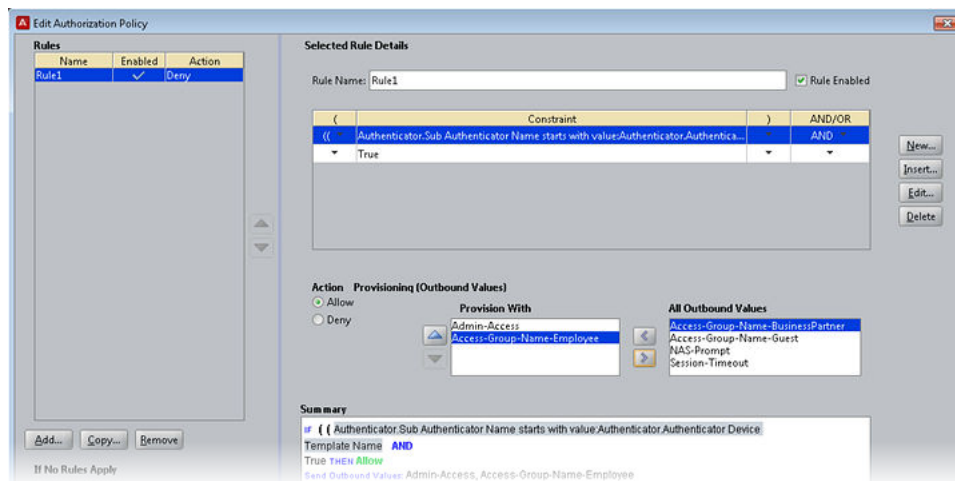
Structure of user Authorization Policies

This section provides an introduction and reference to user authorization policies and the Dashboard windows you use to edit them.

Elements of a User Authorization Policy


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 256).
- 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 274 for details on provisioning instructions.



How Ignition Server evaluates a user Authorization Policy

Ignition Server makes a user authorization decision using a *first-match-wins* logic. The following figure shows an policy example with three rules in the rule sequence



In the preceding 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 255), 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 255), 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 301.)

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 previous 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](#) on page 256.
- **system attributes:** data that describes the date and time of the access request. See [System attributes](#) on page 257.
- **inbound attributes:** values passed by the authenticator in the form of RADIUS attributes or VSAs or from user certificate. These typically describe the context or originating user/device of the access request. See [Inbound Attributes](#) on page 286.
- **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 [Managing Ignition Server licenses](#) on page 75.
- **device attributes:** data that describes the connecting client device such as a user’s laptop or a printer. See [Device Attributes](#) on page 259.

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 229.

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 225.
- **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 [Inbound Attributes](#) on page 286.

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 229.

System attributes

System attributes describe the date and time of the access request. Note that Ignition Server timestamps each incoming transaction with the time of the locale *where Ignition Server is installed*.

System attributes are:

- **Date**, **Date and Time**, and **Time**: These attributes let you use the Ignition Server date and time in a rule. Use the clock icons to set time periods, and use the drop-down list to select the appropriate time zone. See [Using Time and Date in a rule](#) on page 260.

- **False:** always evaluates to false
- **True:** always evaluates to true
- **Weekday:** the weekday range of the Ignition Server, for example, Monday to Friday.

Inbound Attributes

Inbound attributes describe 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.

Inbound Attributes are:

- **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 148. **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 148.
- **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 begin 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 287. If the attribute you want to evaluate does not appear in the list, configure a new inbound attribute as explained in [Preparing an inbound Attribute for use in an Authorization Rule](#) on page 286.
- **User certificate attributes:** Common Name, Country Code, E-mail Address, Locality, Organization, Organization Unit, and State/Province.

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:

- **account-locked:** indicates if the user account is locked.
- **device-address:** MAC address of the device
- **device-compliance:** indicates if this device is compliant with MDM server policies.
- **device-disk-encryption:** indicates if the disk encryption feature on the device is turned on.
- **device-group-member:** indicates the device’s group membership, as recorded in its Ignition Server device record.
- **device-jailbroken:** indicates if the device is jailbroken.
- **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-os-type:** indicates the type of operating system on the device.
- **device-os-version:** indicates the version of operating system on the device.
- **device-pin-lock:** indicates if the pin lock feature on the device is turned on.
- **device-registered:** indicates whether or not the device is registered and active.
- **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-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 [Learned Devices tab](#) on page 473.
- **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 235.

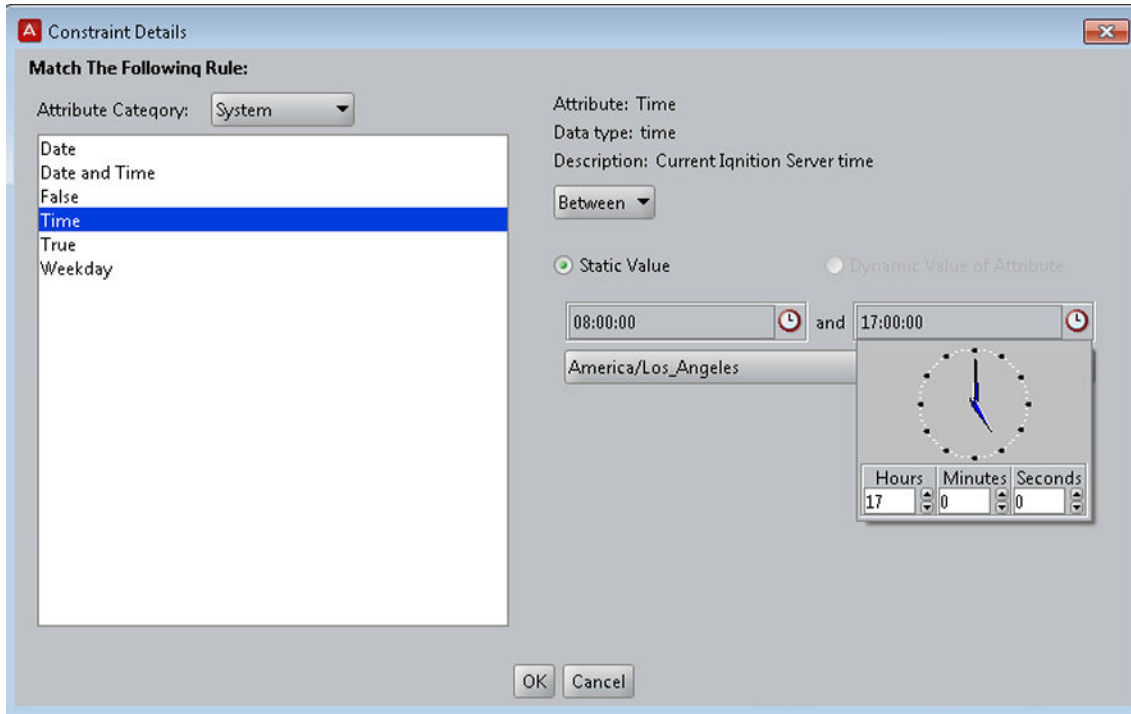
Authenticator attributes

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 218 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 98.
- **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

The system attributes **Date**, **Date and Time**, and **Time** let you use the current Ignition Server date and time 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*.

For example, 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.:

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:

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.

Rules are assembled using:

- **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.

Important:

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.

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 (accountlocked 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.

- **Contains Any:** 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.
- **Contains All:** 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 **Contains Any** comparison operator instead. The **Contains Any** comparison is less strict and might be the correct choice in some cases.
- **Does Not Contain Any:** 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.

- **Does Not Contain All:** 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.

Procedure

1. Select the Access Policy

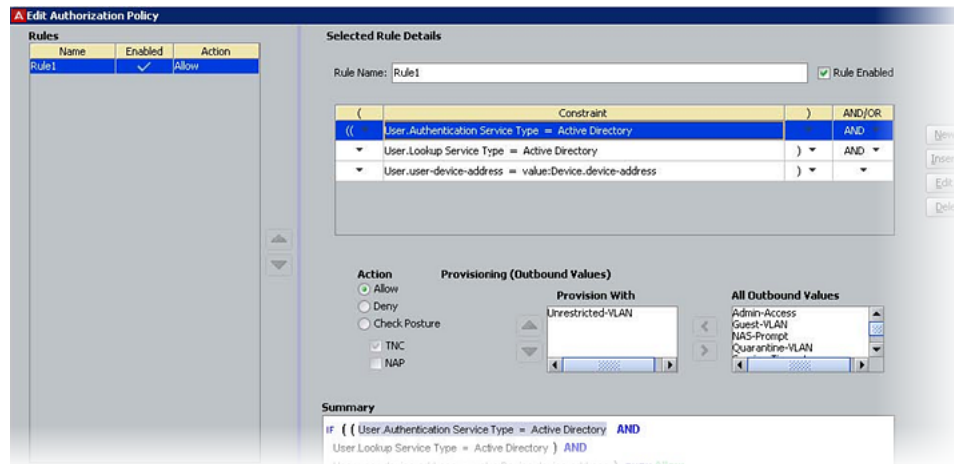
Each user authorization policy applies within the scope of its access policy.

- a. In Dashboard's Configuration hierarchy tree, expand Access Policies and expand RADIUS.

- b. 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

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.



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

- a. 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.

Alternatively, you can copy an existing rule. See [Copying an authorization rule](#) on page 270.

- b. 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

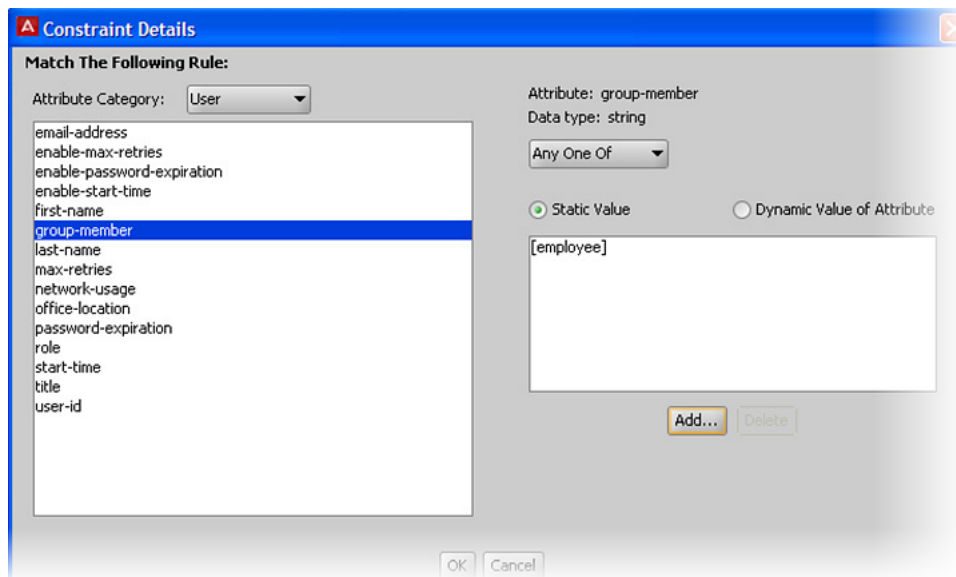
To view a rule, click on its name in the **Rules** list at the left of the Edit Authorization Policy window. The rule appears in the **Selected Rule Details** section of the window. Each rule consists of one or more **Constraints**, an **Action**, and optionally, **Provisioning values**.

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:


- On the left side of the Edit Authorization Policy window, click on the name of the **Rule** you want to edit.
- To the right of the **Constraint** table, click the **New** button. The Constraint Details window appears.



- 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 256.)
- 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.
- 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.
- 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.

- g. Click **OK** to close the Constraint Details window.
- h. 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.
- i. 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.



- j. Use the **AND** and **OR** conjunctions to form a logical condition statement.
 - k. 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 253 for information on the precedence of Allows and Denies in Ignition Server.
 - l. Optional: You may add provisioning to the rule.
 - m. 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.

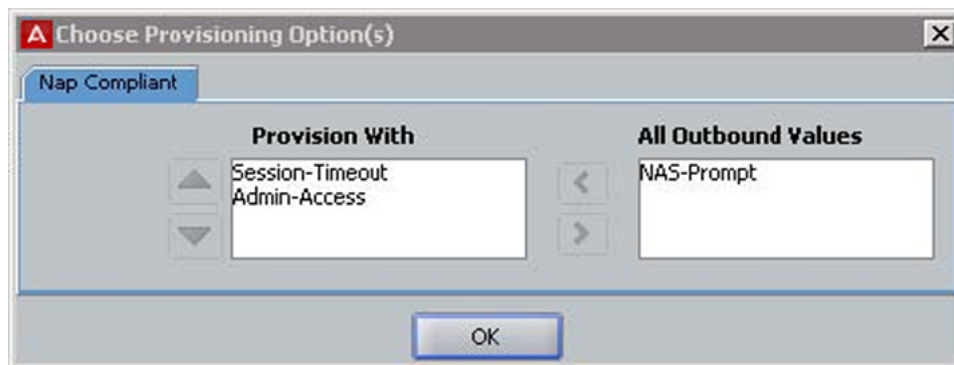
Outbound values sent from an *Allow* rule are typically used to set characteristics of the user's session, while outbound values sent from a *Deny* rule are typically used to convey information about why the denial occurred.

! 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:


- a. Open the **Edit Authorization Policy window**.
- b. Click on the Rule to which you want to add provisioning.
- c. Check **Action** to make sure the *Allow* or *Deny* is configured as desired.
- d. Select the rule you want to provision, and then click **Edit**. The Compliance Condition list displays the NAP rules.
- e. In the **Provisioning (Outbound Values)** section, select the outbound value that 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.



- To see the actual RADIUS attribute name and value to be sent, right-click an outbound value name in either list. A dialog window displays, 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 [Provisioning policy](#) on page 274.

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.

Enabling or disabling rules within a policy

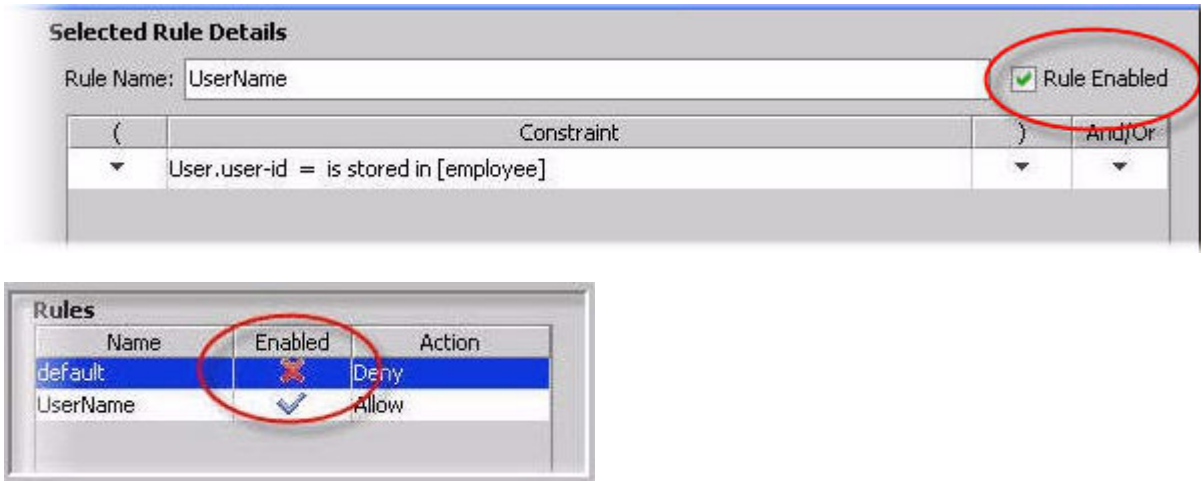
Ignition Server also allows you to enable or disable rules within a user authorization policy. This feature of Ignition Server lets you temporarily activate or deactivate an individual rule, without permanently deleting it from your policy.

Follow this procedure to determine if the rules are enabled.

Procedure

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** check box.
 - If this check box is selected, the highlighted rule is enabled (or active). Selecting 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 clear the **Rule Enabled** check box.

Clearing this check box tells Ignition Server to bypass the highlighted rule when evaluating the **Rules** list of your authorization policy.



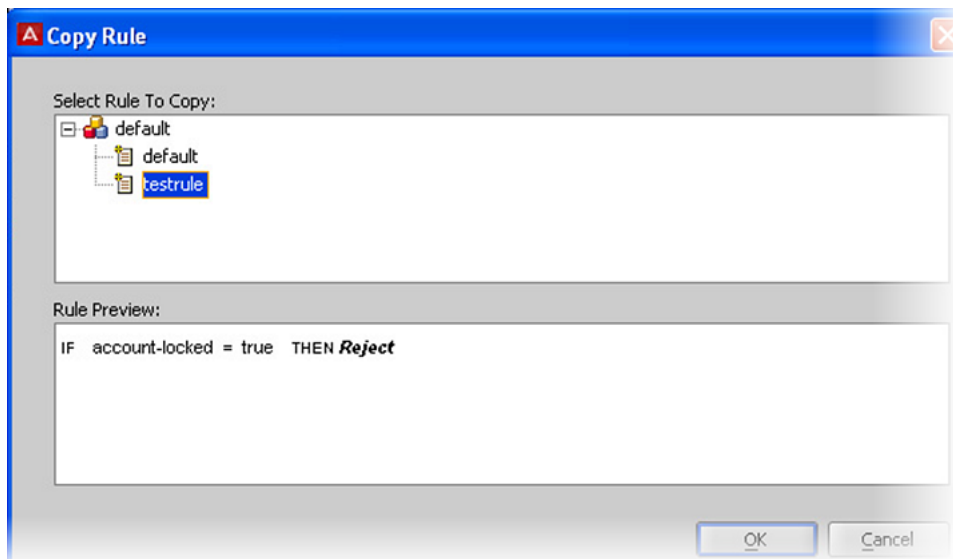
5. Click **OK** to close the window and save the changes to your policy.

Copying an authorization rule

You can copy rules within an access policy or from one access policy to another.

Procedure

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.

Important:

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.

Procedure

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 preceding procedure is to modify the default rule in the “default” Access Policy, making it an authentication-only rule, as follows.

Procedure

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 through 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.

Procedure

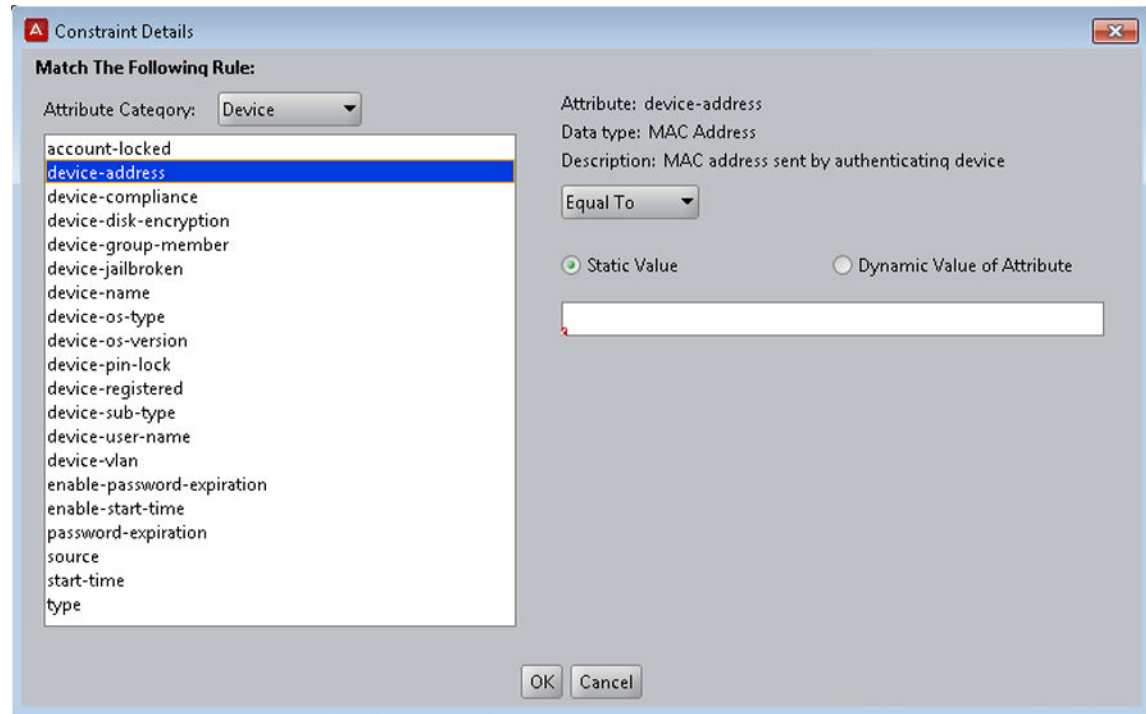
1. Open your authorization policy: In Dashboard’s Configuration hierarchy tree, expand Access Policies and expand RADIUS. Click the name of your access policy. Click the **Authorization** tab and click **Edit**.
2. Write your authorization rule.
 - In the **Authorization Policy** window, create a new rule or edit an existing one.
Click on the rule’s name to edit the rule.
 - In the **Selected Rule Details** section, click **New** to create your constraint.
 - In the **Constraint Details** window, select **Device** from the **Attribute Category** drop-down list.

In the list just below this, click the name of your device attribute. For a complete list, see [Device Attributes](#) on page 259.

If the desired attribute is not there, add it as shown in [Adding Virtual Attributes for Devices](#) on page 235.

On the right side of the window, define the logic of your constraint.

Click **OK**.



- In the **Authorization Policy** window, with your rule still selected, select the desired **Action**. If you want to send provisioning values, go to the **Provisioning** section and select the check box next to each value you want to send.

3. Click **OK** to close the Authorization Policy window.

Chapter 16: Provisioning policy

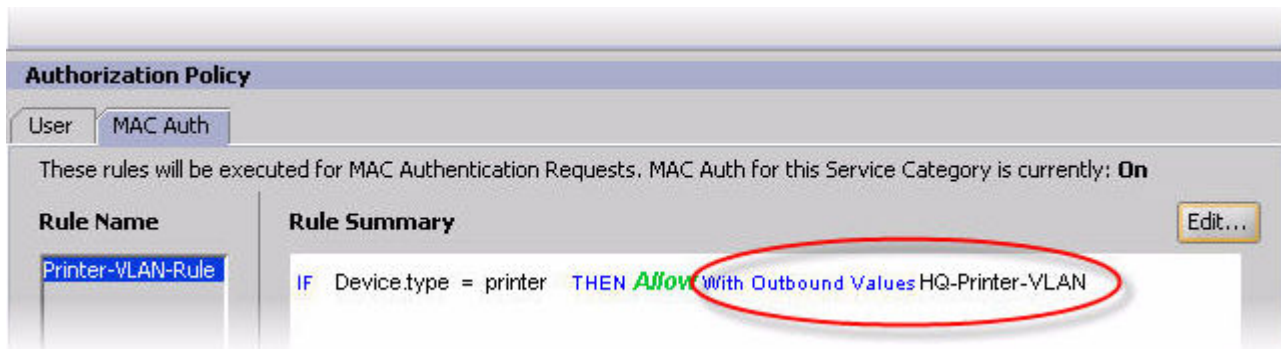
When a user or device authenticates, the Ignition Server policy engine provisions the user's or device's network session by sending instructions that make switch settings, set session time-outs, or assign the user or device to a VLAN. Ignition Server sends these instructions in the form of RADIUS attributes 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 286, this chapter explains how to set up these features of Ignition Server.

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. Configuring 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. In Ignition Dashboard, your provisioning instructions are part of your user authorization and/or MAC authorization rules, as configured through 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.



Setting up session provisioning

Procedure

1. Create the *outbound attribute* as explained in [Outbound Attributes](#) on page 276. 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 value](#) on page 279.
3. Specify the conditions that will trigger Ignition Server to send your outbound value.
 - include your outbound value in a device template and apply that device template to your authenticator definition [Device Templates](#) on page 291; 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 264.

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 281.
- **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](#) on page 292, [Modifying a Device Template](#) on page 294 or [Applying a device template to your authenticator](#) on page 294
- [Listing Ignition Server's set of available RADIUS Attributes](#) on page 295 and [Listing Ignition Server's set of available VSA attributes](#) on page 297.
- [Adding a new RADIUS Attribute](#) on page 296 and [Adding new VSA](#) on page 298.
- [Adding equipment vendor](#) on page 298
- [Overriding the outbound attribute type for one or more authenticators](#) on page 278

- [Finding an Inbound Attribute](#) on page 287 and [Creating a Vendor-Specific Inbound Attribute](#) on page 289.

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](#) on page 276; 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 278.

Finding a global outbound attribute

Procedure

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, use the following procedure.

Procedure

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 one of the following.
 - *To use a standard RADIUS attribute*, select **RADIUS Attribute**, 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 296).

Click **OK**.

- To use a VSA, select **VSA**. In the **Vendor** drop-down list, choose the name of the manufacturer of the authenticator equipment you provision, and in the **VSA** drop-down list, 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 equipment vendor](#) on page 298 or [Adding new VSA](#) on page 298.

5. Your new attribute now appears in the list in the **Outbound Attributes** panel.

Overriding the outbound attribute type for one or more authenticators

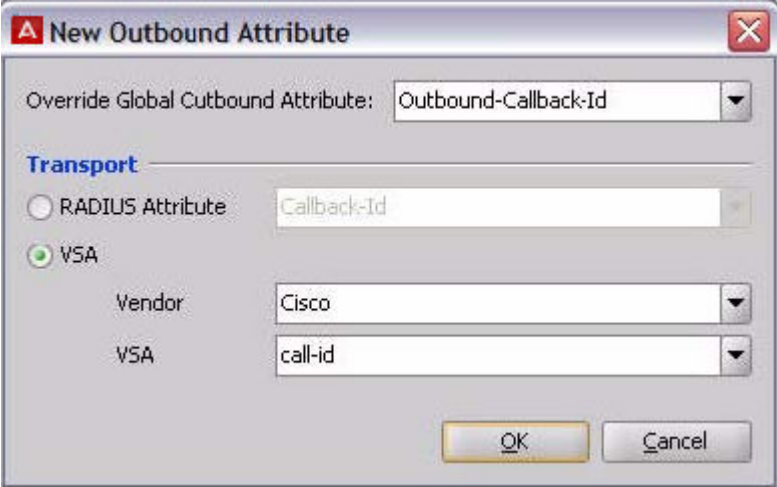
You can create an override which forces Ignition Server to use a *different RADIUS attribute than usual* when sending to a specific authenticator or authenticators. We refer to this as “overriding a global outbound attribute.”

To specify a non-standard outbound attribute to be used in RADIUS messages to authenticators of a single type, follow the steps below.

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 276.

Procedure

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 292.
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.



New Outbound Attribute

Override Global Outbound Attribute: Outbound-Callback-Id

Transport

RADIUS Attribute Callback-Id

VSA

Vendor: Cisco

VSA: call-id

OK Cancel

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 296); 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 new VSA](#) on page 298 or [Adding equipment vendor](#) on page 298).
7. Click **Ok**.

Outbound value

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](#) on page 279.

Finding an outbound value

Procedure

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 279 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.

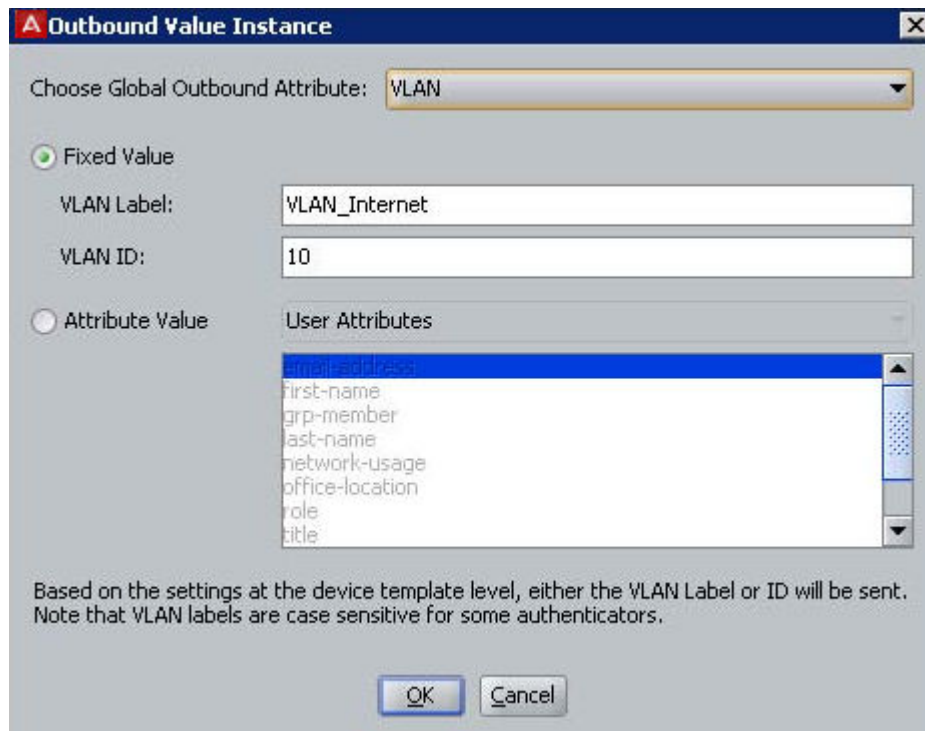
Procedure

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 276.
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.

- 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.

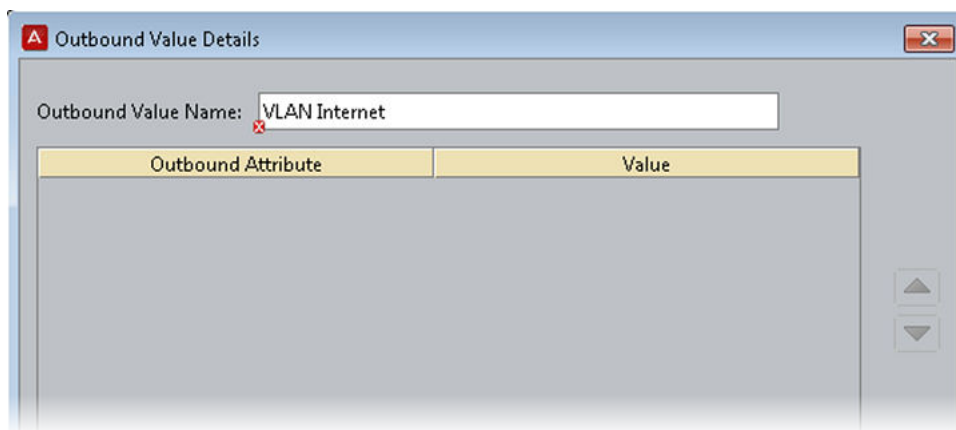


- Click **New** to begin adding a name-value pair that will be sent. Most outbound values send one name-value pair, but you can send as many as needed. The Outbound Value Instance window lets you create each name-value pair.
 - In the **Choose Global Outbound Attribute** drop down list, select the name 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 282.



- 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 267.

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. The built-in outbound values are as follows.

- **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 configure 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](#) on page 282.
- information from the user's record. See [Passing value from the user record or device record to an outbound value](#) on page 283.
- information from the authenticator. See [Passing an inbound value to an outbound value](#) on page 285.

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, use the following procedure.

Procedure

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 that 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 cannot be edited.
5. In the **Value** section, click the upper radio button. (Or, in the VLAN version of the window, click **Fixed Value**.) 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.

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.

Procedure

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 229 or [Device Virtual Attributes](#) on page 234.
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.)

- 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.

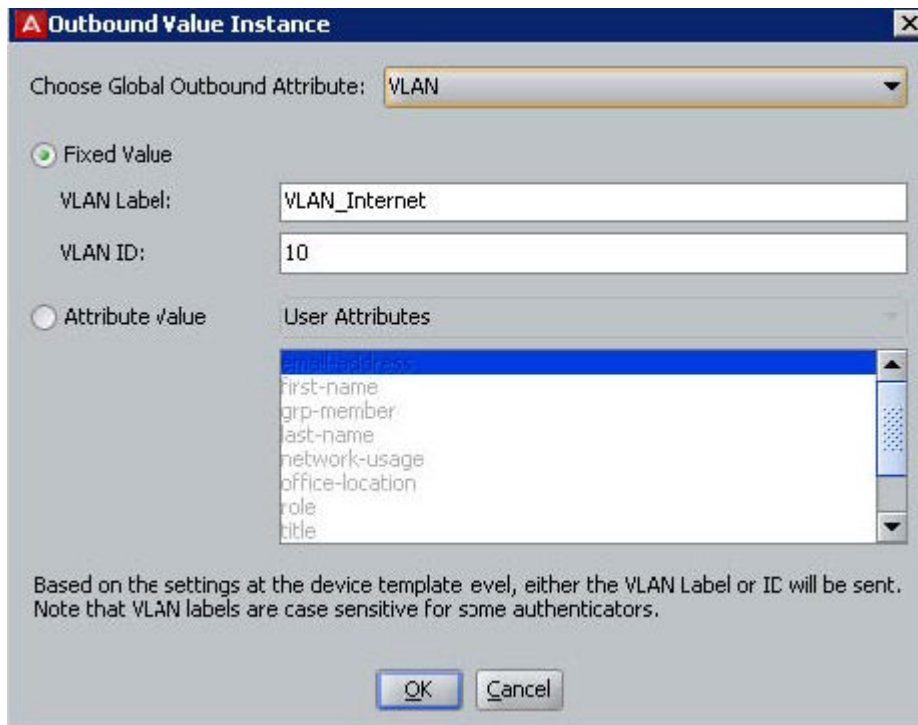
! Important:

Make sure that the outbound attribute you have chosen has the correct data type for the value you want to send.

- In the **Value** section, tick the **Attribute Value** radio button.
- In the drop down list to the right of the radio button, select **User Attributes** or **Device Attributes**.
- 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 230. 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 300.)

The specified virtual attribute's value is retrieved from the user record and placed in the outbound value at runtime.

- Click **Ok**.



- 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.

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.

Procedure

1. For each inbound value that you want to use, define an Ignition Server inbound value as explained in [Inbound Attributes](#) on page 286.
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 300.)

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.

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* or in a user certificate. 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 286; and/or
- return the inbound attribute in the RADIUS response. See [Passing an inbound value to an outbound value](#) on page 285.

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 286.)

Procedure

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 **RADIUS Attribute** 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 296.)
 - If the attribute is a VSA, click **VSA**, 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 equipment vendor](#) on page 298 or [Adding new VSA](#) on page 298.)
5. Click **Ok**.
6. Define your authorization rule to evaluate the inbound attribute.

- a. 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**.
- b. Create your rule by clicking **New (under the Rules list)**, giving the rule a name, and clicking **OK**. (You can also edit an existing rule).
- c. In the **Selected Rule Details** section, click **New** to add a constraint.
- d. In the **Constraint Details** window, in the **Attribute Category** drop-down list, select **Inbound**.
- e. In the list of inbound attributes, select the inbound attribute you saved earlier.
- f. On the right side of the window, define the comparison condition that must be met in order to trigger this rule. (For more information, see [Creating a RADIUS user authorization policy](#) on page 264).
- g. Click **OK**.

Finding an Inbound Attribute

Follow this procedure to find an inbound attribute.

Procedure

1. View the list of available *global* inbound attributes by expanding the **Provisioning** node of Dashboard's **Configuration** tree and clicking **Inbound Attributes**.



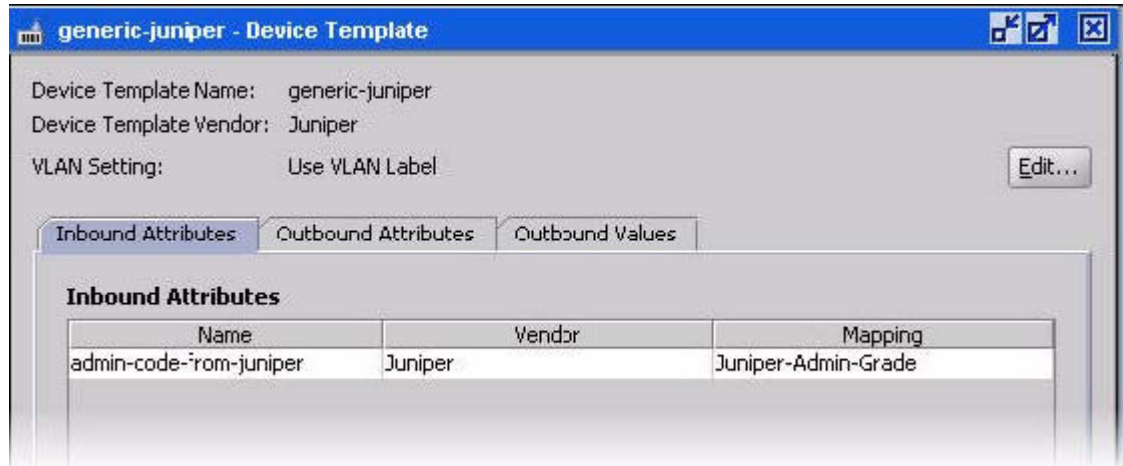
Name	Vendor	Mapping
Admin-Access-Request	RADIUS	Service-Type
Inbound-Callback-Number	RADIUS	Callback-Number
Inbound-Called-Station-Id	RADIUS	Called-Station-Id
Inbound-Calling-Station-Id	RADIUS	Calling-Station-Id

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 previously.



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.

Procedure

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).

The screenshot shows a dialog box titled "New Inbound Attribute". It has a text field for "Inbound Attribute" containing "juniper-user". Below this is a section titled "Transport" with two radio buttons: "RADIUS Attribute" (unselected) and "VSA" (selected). The "RADIUS Attribute" radio button is associated with a dropdown menu showing "Acct-Authentic". The "VSA" radio button is associated with two dropdown menus: "Vendor" showing "Juniper" and "VSA" showing "Juniper-Local-User-Name". At the bottom are "OK" and "Cancel" buttons.

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 296).
 - 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 new VSA](#) on page 298 or [Adding equipment vendor](#) on page 298.
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 286) or you return the inbound value in the RADIUS response (see [Passing an inbound value to an outbound value](#) on page 285).

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.

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 288.

Define the vendor-specific inbound attribute as follows:

Procedure

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 292.
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 296).
 - 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 new VSA](#) on page 298 or [Adding equipment vendor](#) on page 298).
7. Click **OK**.

Next steps

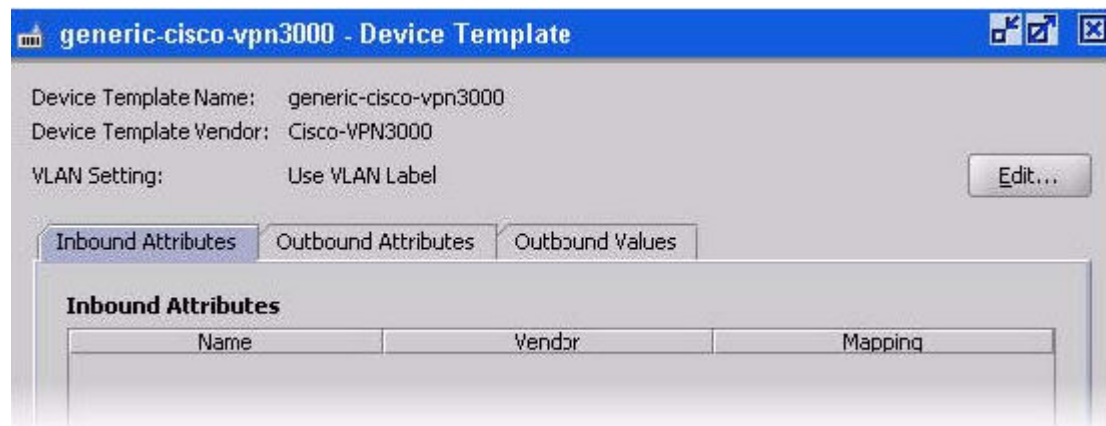
Now that you have finished creating the inbound attribute, you can evaluate its value in an authorization rule (see [Inbound Attributes](#) on page 286), or you can return the attribute value in the RADIUS response to the authenticator (see [Passing an inbound value to an outbound value](#) on page 285).

Device Templates

Ignition Server has a configuration tool called a *device template* 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 218), 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).

Finding a device template

Use the following steps to find a device template in Ignition Server.

Procedure

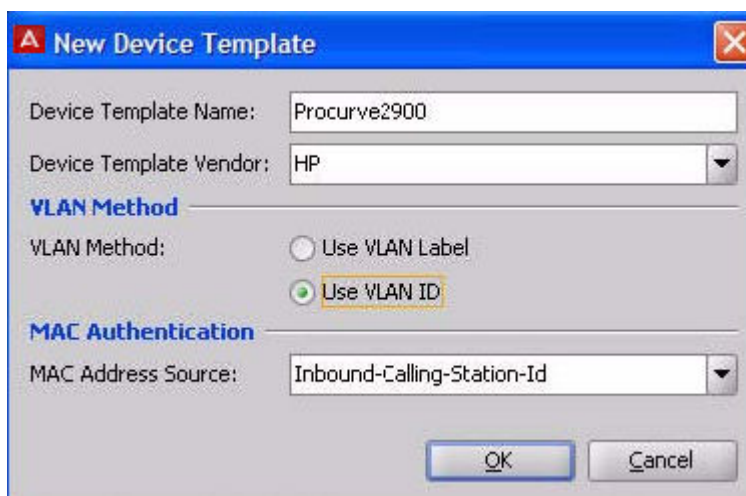
1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Vendors/VSAs**.
2. In the left panel of the **Vendors** panel, double-click the manufacturer name of your network equipment, then click **Device Templates** to display the list of templates for that manufacturer.

Creating a Device Template

Use the following steps to create a device template:

Procedure

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 equipment vendor](#) on page 298).
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.

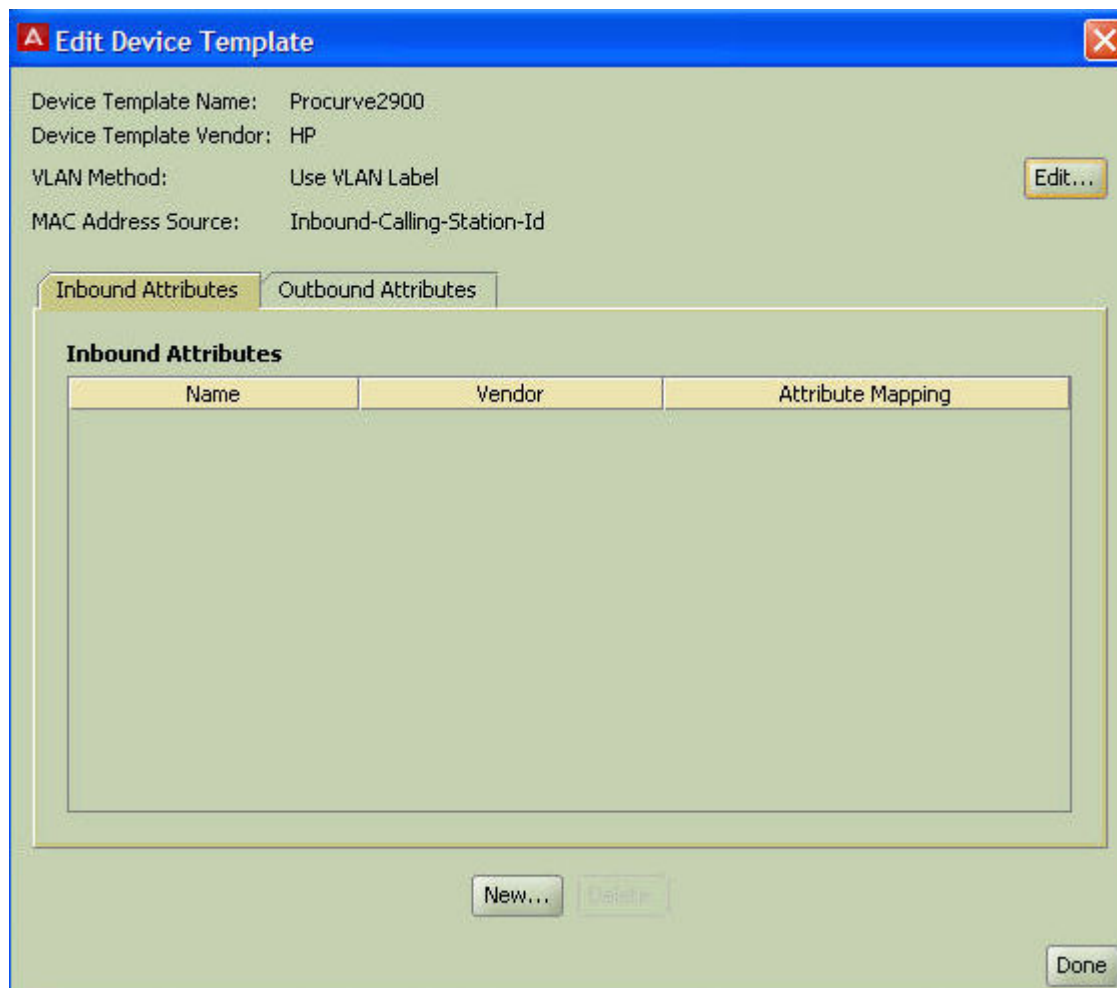


- In the **MAC Address Source field**, choose the RADIUS Attribute that contains the MAC address of connecting devices.

⚠ 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.

- 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.



- Apply your device template to each authenticator that is to use it. See [Applying a device template to your authenticator](#) on page 294.

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 286.
- 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 278.

Modifying a Device Template

To change a device template, use the following steps.

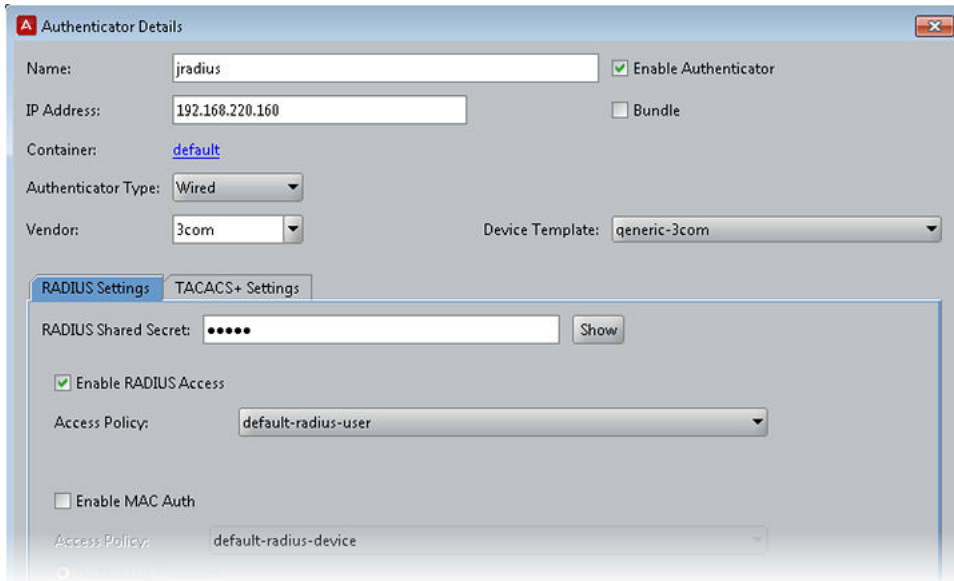
Procedure

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 286.
 - To set up **Outbound Attributes** and corresponding **Values**, see [Creating a global outbound attribute](#) on page 276.

Applying a device template to your authenticator


Procedure

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 Procedure

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.

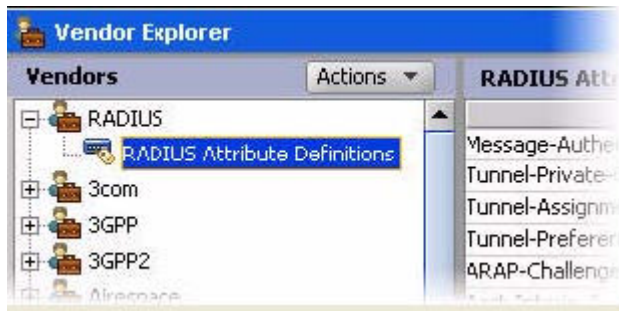
To add a new RADIUS attribute, see [Adding a new RADIUS Attribute](#) on page 296.

Adding a new RADIUS Attribute

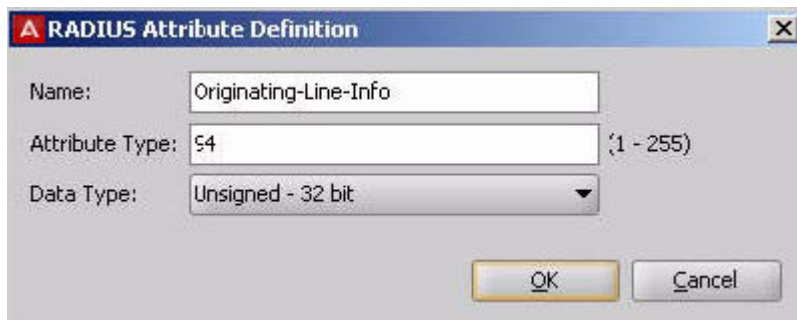
If the RADIUS attribute you want to use does not appear in Ignition Server's default list of RADIUS attributes (you can view this list in the Vendors panel), create it using these steps. Note that the following steps apply only to standard RADIUS attributes. If you want to create a new vendor-specific attribute, see [Adding new VSA](#) on page 298 instead.

Procedure

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**.
5. 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 276.
- To evaluate the value of this RADIUS attribute in your authorization rules, turn next to [Inbound Attributes](#) on page 286.

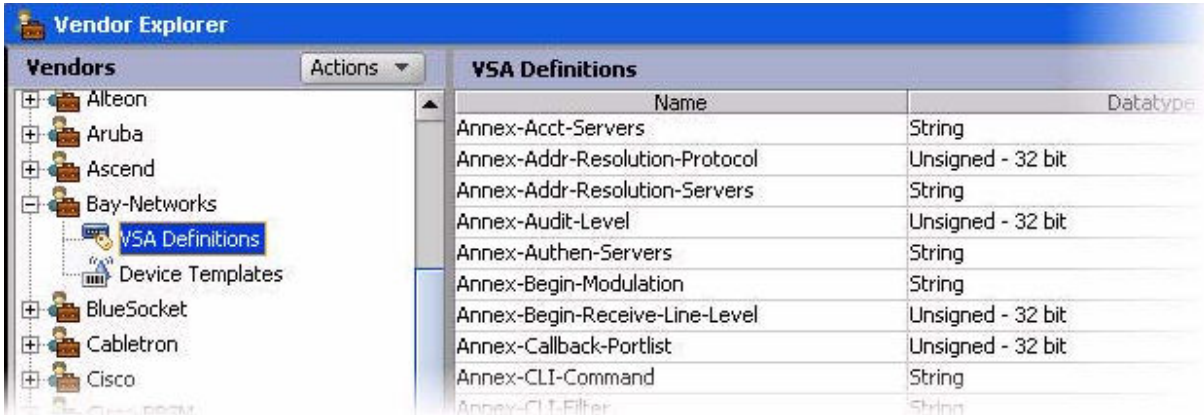
Listing Ignition Server's set of available VSA attributes

Follow this procedure to view a list of the vendor-specific RADIUS attributes (“VSAs”) defined in Ignition Server.

Procedure

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 screenshot shows the 'Vendor Explorer' window. On the left, a tree view lists vendors: Alteon, Aruba, Ascend, Bay-Networks, BlueSocket, Cabletron, and Cisco. Under 'Bay-Networks', 'VSA Definitions' is selected. On the right, a table titled 'VSA Definitions' displays the following data:

Name	Datatype
Annex-Acct-Servers	String
Annex-Addr-Resolution-Protocol	Unsigned - 32 bit
Annex-Addr-Resolution-Servers	String
Annex-Audit-Level	Unsigned - 32 bit
Annex-Authen-Servers	String
Annex-Begin-Modulation	String
Annex-Begin-Receive-Line-Level	Unsigned - 32 bit
Annex-Callback-Portlist	Unsigned - 32 bit
Annex-CLI-Command	String
Annex-CLI-Filter	String

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.

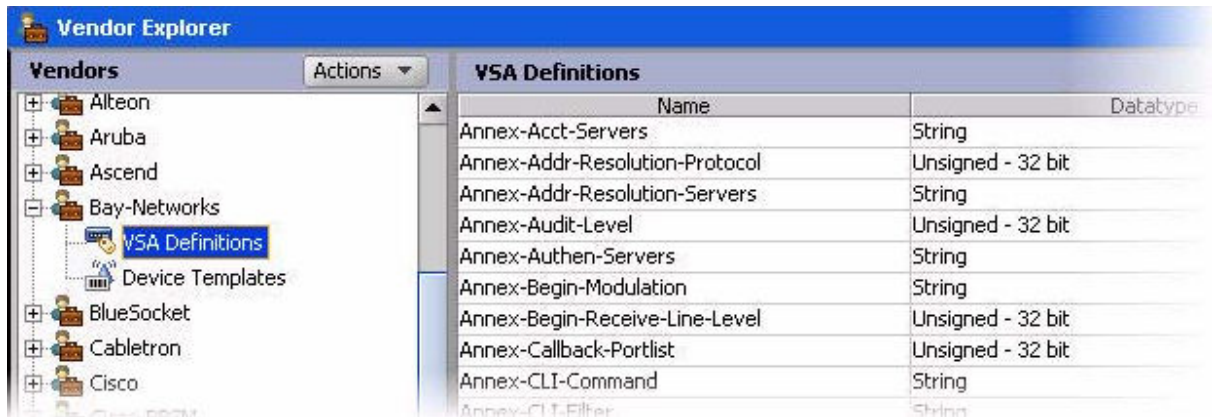
The **Attribute Type** is the integer code that designates the attribute, as specified in your network equipment's documentation, or in the relevant industry standards document.

Adding 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.

Procedure

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 equipment vendor](#) on page 298.



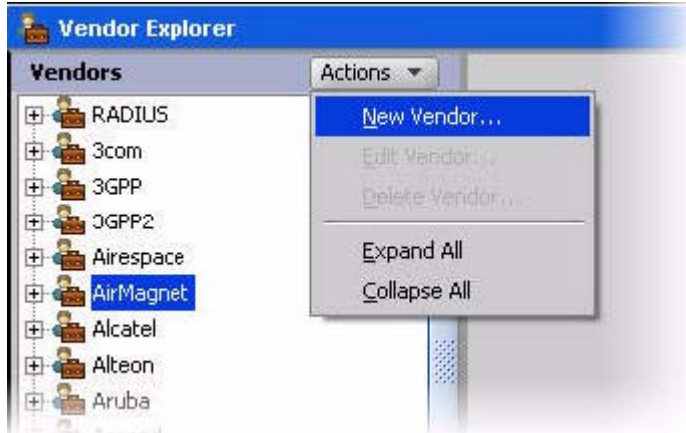
3. At the bottom of the Vendors panel window, click **New**.
4. In the RADIUS VSA Definition window, define the attribute.
 - Enter the **RADIUS VSA Name** without spaces. This is the RADIUS attribute name and must match the name used by your networking equipment.
 - Enter its **Attribute Type** as an 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**.
5. Click **OK**.

Adding 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.

Procedure

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 298.

To create a device template for the equipment, see [Device Templates](#) on page 291.



Provisioning FAQ

Question	Answer
<p>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?</p>	<p>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 278.</p>
<p>How can I find out the datatype of an inbound or outbound attribute?</p>	<p>Perform the following steps:</p> <ol style="list-style-type: none"> 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. 3. In Dashboard's Configuration tree, expand the Provisioning node and click Vendors/VSAs. 4. Perform the following steps. <ul style="list-style-type: none"> • 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. 5. Scroll to find your attribute. The Data Type column shows the datatype.

Chapter 17: 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 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 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 noncompliant; the policy will have options to determine a course of action depending on 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.

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 253), 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

Enable NAP services on the client

Procedure

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

Procedure

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.

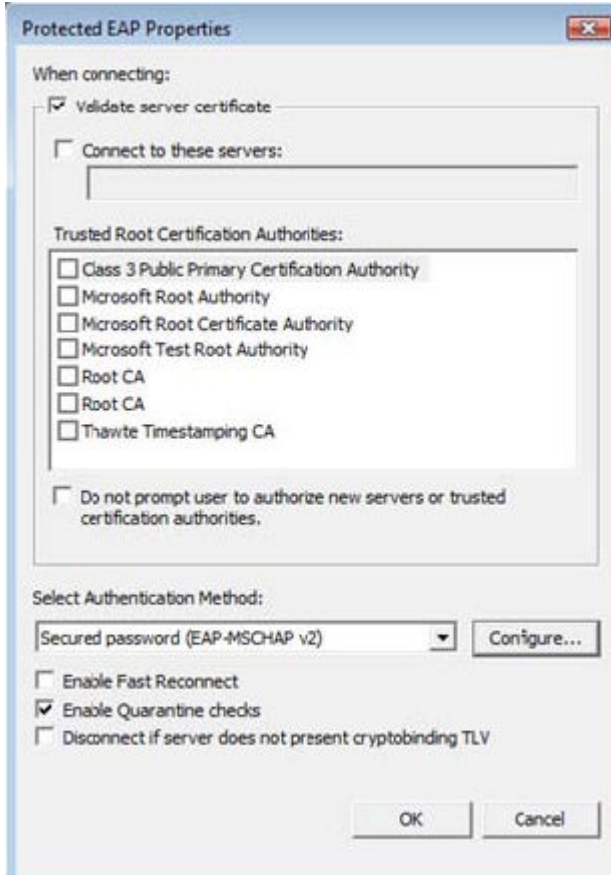
Procedure

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.

Procedure

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**.

Edit Posture Profile

Posture Profile Name: Posture_Profile

NAP Configuration

Posture

Product	Enabled	Up to date	Comment
Anti-Virus	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	An antivirus application is active and up to date
Anti-Spyware	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	An antispymware application is active and up to date Note: This option is "Not" applicable for Windows XP client
Firewall	<input checked="" type="checkbox"/>		A firewall application is enabled for all network connections
Windows Automatic Update	<input type="checkbox"/>		Automatic updating is not enabled

Windows Security Update Protection

Restrict access for clients that do not have all available security updates installed

Specify the minimum security level required for updates: All

Maximum number of hours allowed since client has checked for new security updates: 72

By default, clients can receive security updates from Microsoft update. If additional sources are required for deployment, select one or both of the following sources.

Windows server update services

Windows updates

Remediation

Probation time: 2014-02-27 07:42:34
(Clients will be allowed network access till above specified time even if they are found as non-compliant).

URL for remediation server: www.coreavaya.com

Auto remediate

OK Cancel

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 updates to run on the selected product.
Up to date is not applicable to Firewall or Windows Automatic Updates.
4. In the Windows Security Updates Protection, select the **Restrict access for clients that do not have all available security updates installed** check box. Then specify the minimum level of security required for updates.
 - Unspecified (Default)
 - Low and above
 - Moderate and above
 - Important and above
 - Critical only
5. Select the **Windows server updates services** and **Windows updates** check boxes.

6. In the **Remediation** section, select the probation time for when you want NAP to connect to the remediation server for patches and updates. This field can not be edited manually. Click the time/calendar icon to change the hour value.
 - Minimum: 1 hour
 - Maximum: 72 hours (Default)
7. Enter the remediation server URL.
8. Select **Auto remediate** if you want automatic updates.

Chapter 18: VLAN Assignment

This chapter explains how to configure 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 [Provisioning policy](#) on page 274.

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 [Provisioning policy](#) on page 274.

Creating a policy that assigns users to VLANs

Setting up VLAN provisioning in Ignition Server requires two or three steps.

Procedure

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-PrivateGroup-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 309.

- If your authenticator requires that you use a different RADIUS attribute for VLAN assignments, follow the instructions in [Setting Up VLAN provisioning using nonstandard RADIUS Attributes](#) on page 308.

Setting Up VLAN provisioning using nonstandard RADIUS Attributes

Use the following steps to configure VLAN provisioning for VLAN equipment that *does not* accept assignments through 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 309).

Procedure

1. In Dashboard's **Configuration** tree, expand the **Provisioning** node and click **Outbound Attributes**.
2. In the **Outbound Attributes** panel, click **New**.
This step creates a new outbound attribute.
3. 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, [Adding equipment vendor](#) on page 298.
- If the **VSA** list does not include the VLAN attribute name required by your authenticator, see [Adding new VSA](#) on page 298.

5. Click **OK** to save the outbound attribute.

Create an Outbound value for each VLAN

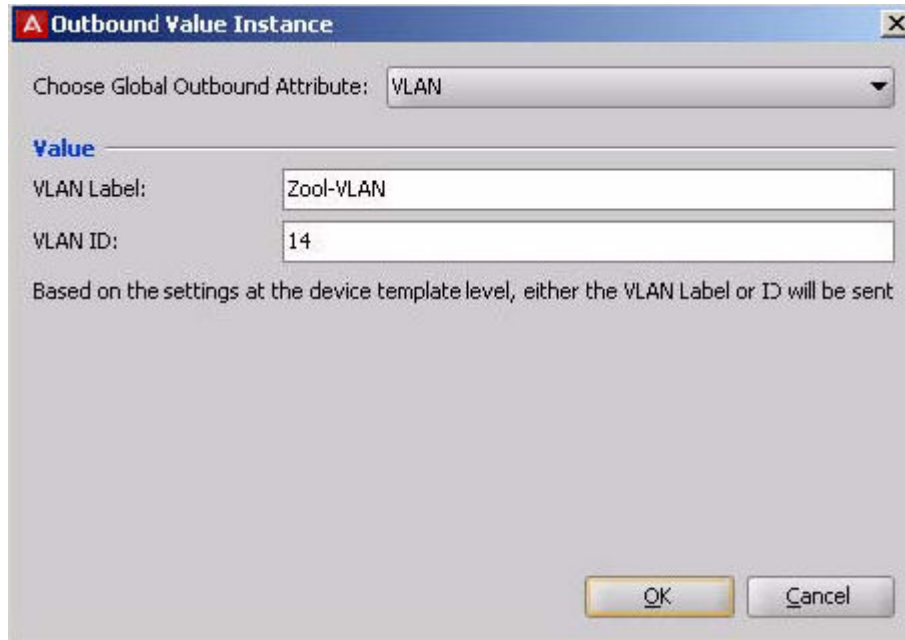
Before you can write a VLAN provisioning rule, you must save the VLAN ID or name as an Ignition Server *outbound value*. The outbound value is sent in the RADIUS message to the authenticator. Follow the steps below to create an outbound value:

Before you begin

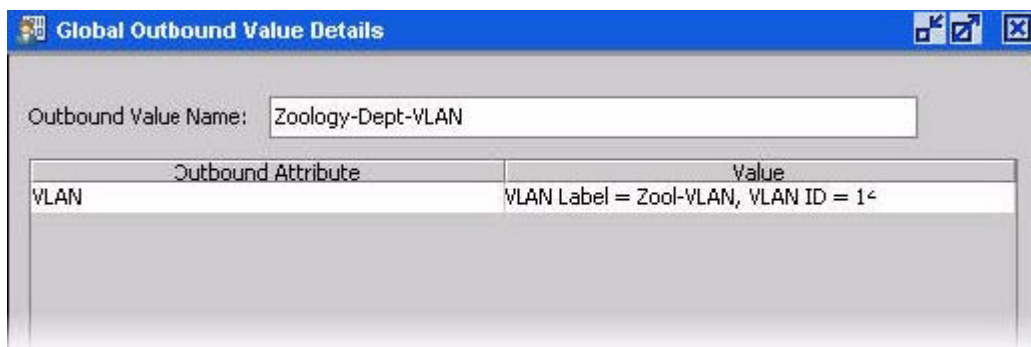
Before you start creating outbound values, log into your VLAN switch as administrator and make a note of the VLAN label (a string) and VLAN ID (an integer) of each VLAN to which you plan to assign users.

Procedure


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 308.
 - 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.

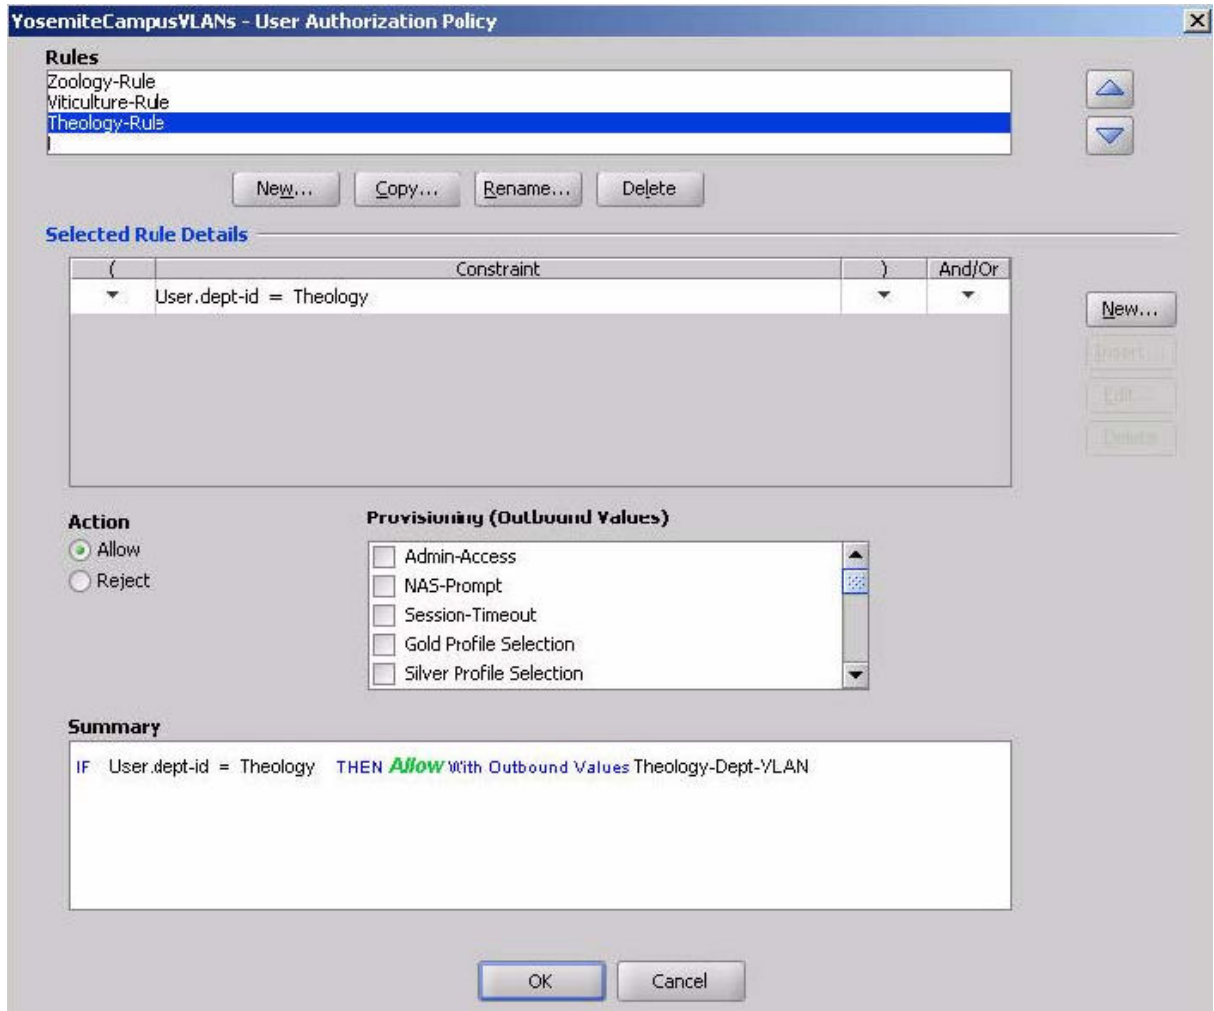
Create VLAN provisioning rules

At runtime, Ignition Server evaluates its set of *provisioning rules* in order to determine to which VLAN it assigns the user or device.

Use the following procedure 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 [Introduction to MAC Authentication](#) on page 339.)

Procedure

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."
3. Click the name of your access policy. Click the **Authorization** tab and click **Edit**.
4. 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 **Edit**.
5. 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.

6. On the left side of the **Edit Authorization Policy** window, click **Add** just below the **Rules** list.
7. In the **New Rule** window, type a name for your provisioning rule and click **OK**.

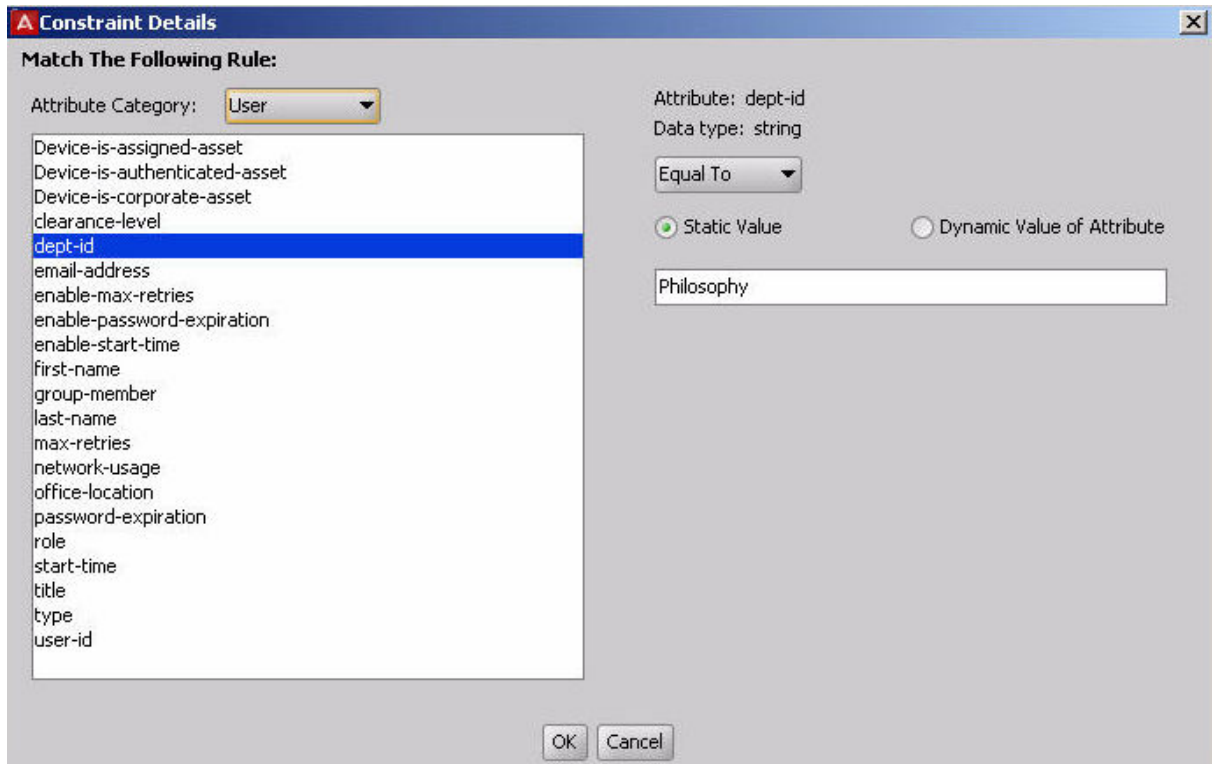


Your new rule now appears in the Rules list in the Edit Authorization Policy window

8. . Click on your new rule 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.

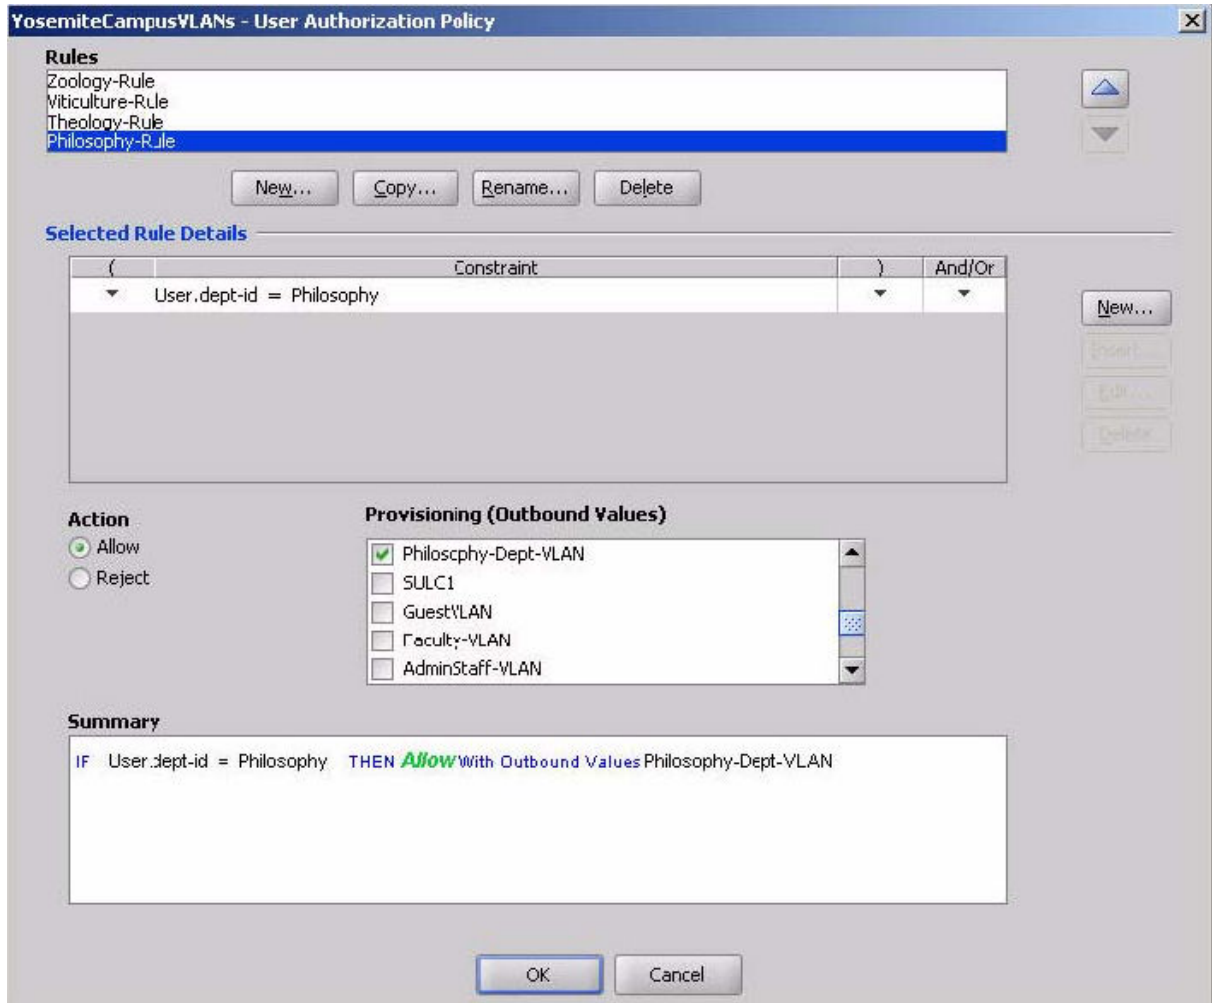
9. Click **New** next to the Constraint list.
10. 9. The **Constraint Details** window defines the condition that must be fulfilled in order to provision the user. Typically, your constraint evaluates a user attribute (see [Attributes used in Rule Constraints](#) on page 256). For this example, we choose a **User Attribute** called “dept-id” and set the test value to “Philosophy.”
 - 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 231).
 - Select **Equal To** .
 - Select the **Static Value** check box.
 - 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.



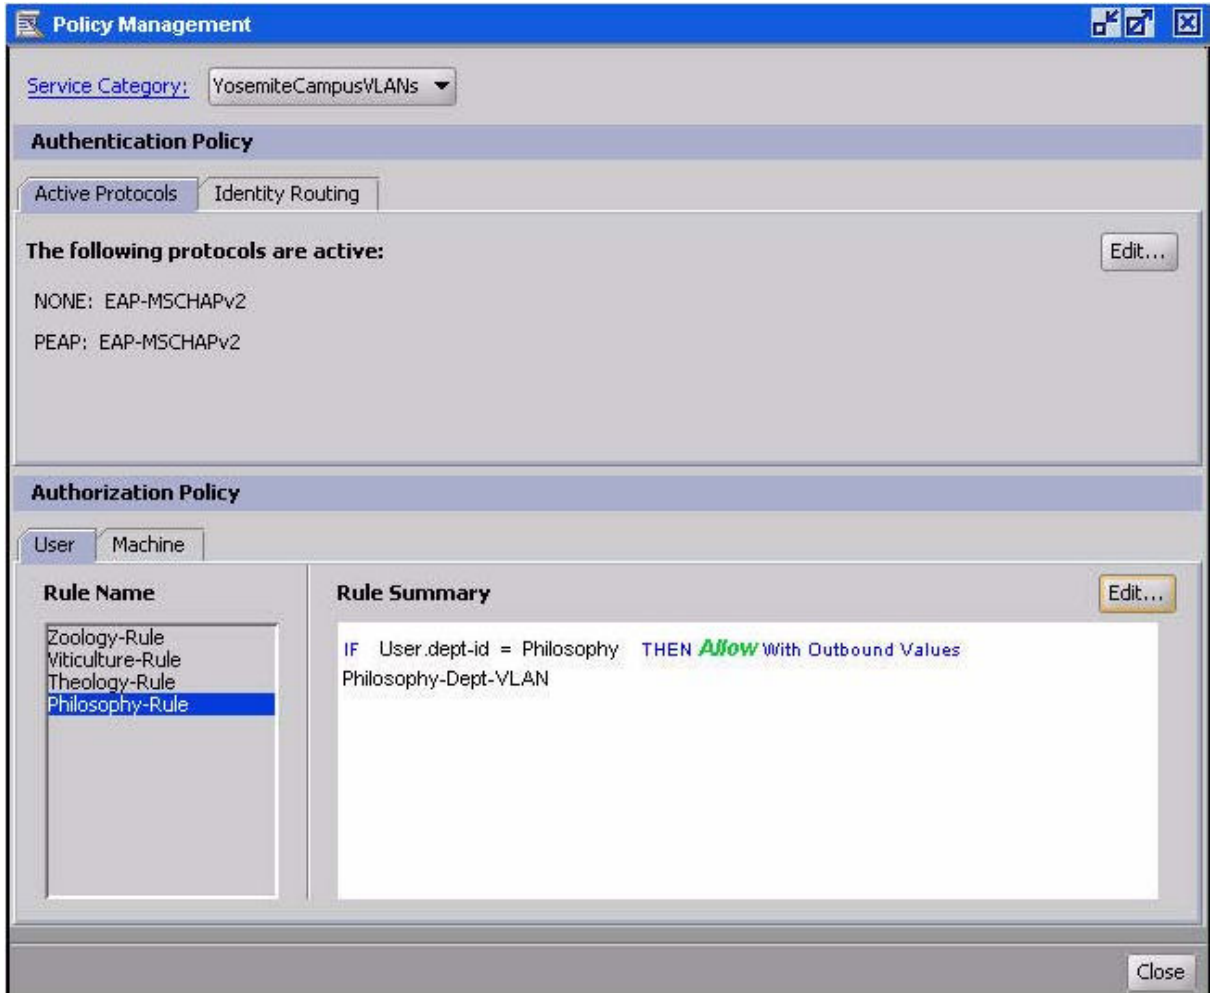
11. In the **Edit Authorization Policy** window, with your rule (“Philosophy-Rule” in this example) highlighted, click **Allow** in the **Action** section, and in the **Provisioning** section select the check box next to the outbound value for the appropriate VLAN.

In this example, we use the outbound value we created in Step 6 in [Create an Outbound value for each VLAN](#) on page 309.



12. Repeat Step 5 through Step 10 to add more VLAN provisioning rules.
13. 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.



Chapter 19: 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.

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 473.

Do not confuse Microsoft Windows machine authentication with the more generic *MAC authentication*. (See [MAC Authentication](#) on page 339.)

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 configure a rule in Ignition Server that only permits the user to log in on a computer that has successfully performed Windows machine authentication. See [Introduction to Asset Correlation](#) on page 353.
 - 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.
- 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.

As the 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 through 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 [MSSOH], 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 319
2. [Machine authentication based on OU, O, or group membership](#) on page 322

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.

Procedure

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

Follow this procedure to create a user virtual attribute called "type" and map it to the AD attribute, objectClass.

Procedure

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, select the name of your AD service.
6. Make sure **Available attribute name** is selected. 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.

Procedure


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 **Add** 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**”.
 - Set **Format** to “Ignore Case”.
 - Select **Static Value**.
 - Click **Add**.
 - In the **Add Value** dialog, type “**computer**” and click **OK**.
This is your comparison value.
 - Click **OK** to close the Constraint Details window.



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 OUs or Os, 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 358.

If your policy needs rules to handle user authentication, return to the top of the window and click **New** to create another rule. See [Creating a RADIUS user authorization policy](#) on page 264.

Set up your supplicants

Configure your supplicants to require machine authentication. Consult your supplicant documentation for instructions. For Microsoft Windows XP supplicants, follow this procedure.

Procedure

1. Open the **Network Connections** window, and open the **Properties** window for the Interface you want to configure.
2. Click **Properties** to open the Properties window, and click the **Authentication** tab.
3. Select the “**Authenticate as computer...**” check box
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 configure 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 319.)

To do this, place your device entries in an OU, O, or group in AD, and configure Ignition Server to allow access based on membership in that group. The following sections provide information on how to perform that configuration.

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 procedures 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.

Procedure

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**.

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.

Procedure

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 "okdevices".
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

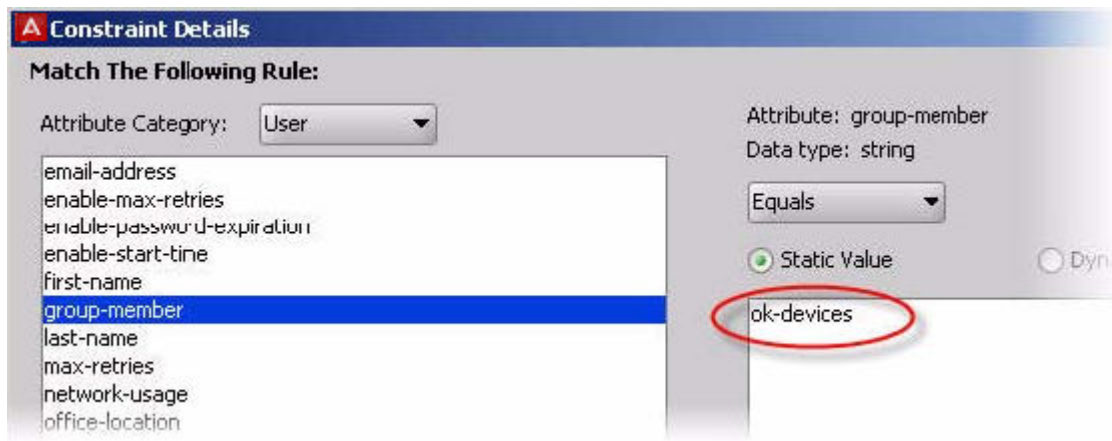
Write a policy rule that checks if the *type* attribute is set to "Computer" and, if so, carries out its machine authentication policy.

Procedure

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.
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, 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. 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 358.

If your policy needs rules to handle user authentication, return to the top of the window and click **New** to create another rule. See [Creating a RADIUS user authorization policy](#) on page 264.

Set up your supplicants

Configure your supplicants to require machine authentication. Consult your supplicant documentation for instructions. For Microsoft Windows XP supplicants, follow this procedure.

Procedure

1. Open the **Network Connections** window, and open the **Properties** window for the Interface you want to configure.
2. Click **Properties** to open the Properties window, and click the **Authentication** tab.
3. Select the “**Authenticate as computer...**” check box
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 through Windows machine authentication. If you have imposed an asset correlation policy (see [Requiring the user to connect using a Machine Authenticated-Device](#) on page 358), 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 473. To configure the TTL, use the following procedure.

Procedure

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.

Chapter 20: TACACS+ authorization

This chapter shows how to configure Avaya Identity Engines Ignition Server as your TACACS+ administrator access control server.

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 327.
- 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 328.

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:

- 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 perform first-time set-up of TACACS+ on your Ignition Server, see [Installing your TACACS+ license](#) on page 329.

- 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 333.
 - if your TACACS+ policy uses *per-command authorization*, create your sets of allowed commands as shown in [Creating a Command Set](#) on page 331.

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 78.

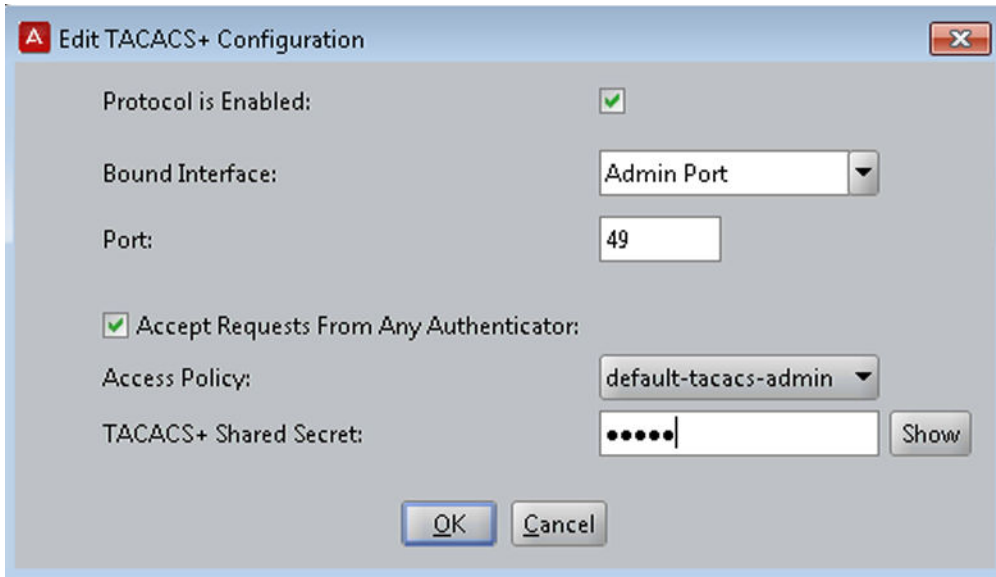
Turning on the Ignition Server TACACS+ service

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 405). Use the TACACS+ tab to bind the TACACS+ service and set its port numbers.

Procedure

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 displays.



4. Edit as necessary:

- **Protocol Is Enabled:** Select this check box 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 Servers, 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 405 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 check box selected, 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:** Select this check box 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 [Using the TACACS+ global authenticator](#) on page 337.
- **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 333.
- If your TACACS+ policy uses *per-command authorization*, create your sets of allowed commands as shown in [Creating a Command Set](#) on page 331.

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.

Procedure

1. In Dashboard's **Configuration** hierarchy tree, expand **Access Policies**, expand **TACACS+**, and click **Device Command Sets**.




Current Site: Sunnyvale Campus	
Device Command Sets	
Name	Description of Set
all-commands	Allows any command
default-command-set	The default command set

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.
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, select the **Simple Command** check box, 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, select the **Allow** check box. To allow only the specific command and arguments you have typed, select the **Deny** check box. Click **OK** to add the command to the list.



OR

- Click **Add** and, in the **New/Edit Device Command Set** window, select the **Advanced Command Matching Regular Expression** check box. 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, select the **Match Command Plus Additional Keywords/Arguments** check box. To allow only the specific command and arguments you have typed, select the **Exact Match for Each Command Only** check box.
 - Click **OK** to import the list.
5. Click **Add** or **Import** again to add more expressions, or click **OK** to save the set.

Viewing or editing a command set

Follow this procedure to view or edit a command set.

Procedure

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 and then specifying the TACACS+ policy name in the authenticator record. (See [Creating an authenticator](#) on page 104.)

Before you begin

Note the following prerequisites, based on the type of authorization you use.

- If you are configuring *per-command authorization*, then you should have already created your command set(s) of allowed commands and arguments. If you have not done this, go to [Creating a Command Set](#) on page 331.
- If you are configuring *privilege-level authorization*, then you should have already assigned a privilege level to each command on the equipment your administrators will manage.


Procedure

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**.

Configure your user lookup policy. See [Creating an Identity Routing policy](#) on page 248.

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 256.)
 - 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.

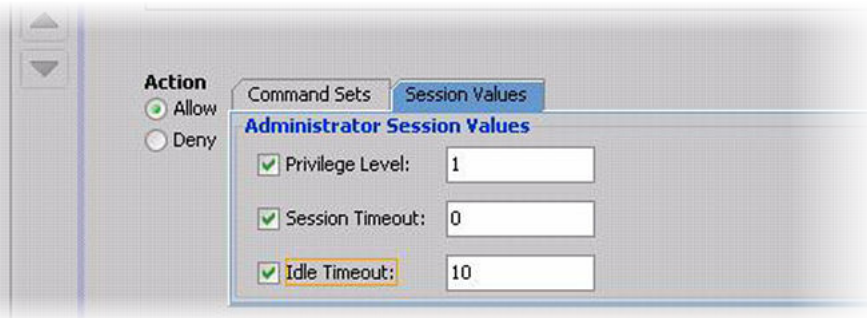
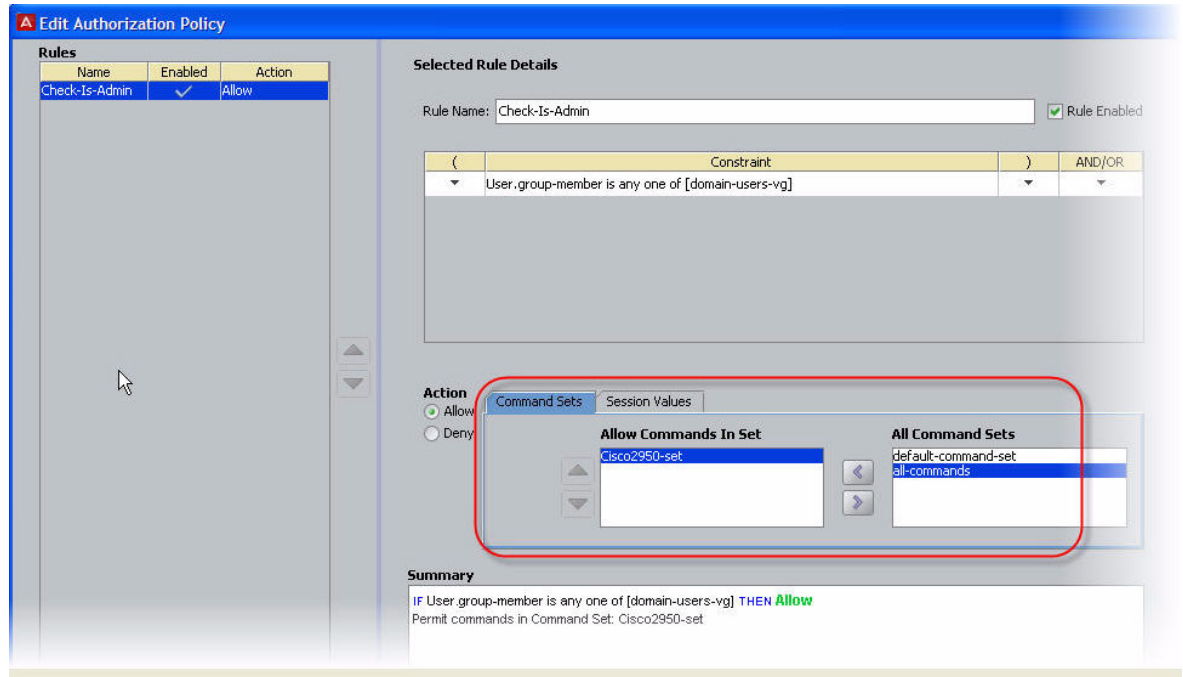
- d. Provide the comparison value by doing one of the following:
 - If you want to compare the attribute value with a fixed test value, select **Static Value** and type or choose the comparison value in the field below that.
 - If you want to compare the attribute value with a value retrieved from another attribute, select **Dynamic Value of Attribute**. 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 **New** or **Insert** 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.
 - 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 configuring *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 configuring Privilege-level authorization, click the **Session Values** tab. Select the **Privilege Level** and enter value of **1-15**. Similarly, select the **Session Timeout** (maximum length of session, regardless of activity) and **Idle Timeout** (maximum amount of time the session can sit idle between commands) in minutes. Time-outs are provisioned by Ignition Server and enforced by the TACACS+ client. A time-out value of zero means that no timeout is entered.



13. Click **OK** to save the policy.

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:

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.

Procedure

1. Configure your device to use Ignition Server as its TACACS+ server. Use your device's command line interface or other tool to configure the following.

- Configure the *TACACS+ server address* with the IP address of the Ignition Server TACACS+ service.

To determine 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.

- Configure *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.

2. In Ignition Dashboard, open the Authenticator Details window as follows: In the **Configuration** hierarchy tree, expand **Authenticators**.

- 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 more information, see [Creating an authenticator](#) on page 104.

3. In the **Authenticator Details** window, click the **TACACS+ Settings** tab.

4. Select the **Enable TACACS+ Access** check box.
5. Type the **TACACS+ Shared Secret** that you specified in Step 1.
6. In the **Access Policy** drop-down list, choose your policy. This is the policy you created in [Creating a TACACS+ Access Policy](#) on page 333.
7. Click **OK** to save the definition.


Using the TACACS+ global authenticator

As explained in [Introduction to Authenticators](#) on page 97, 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.”

Procedure

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, select the **Accept Requests from Any Authenticator** check box.
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**.

Chapter 21: MAC Authentication

This chapter 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.

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 through 802.1X, you can configure 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.

Do not confuse MAC authentication with *Windows machine authentication* and *asset correlation*, which uses Windows machine authentication. (See [Introduction to Windows Machine authentication](#) on page 316.)

Warning:

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 through 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, 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**, place all MACauthenticated 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.

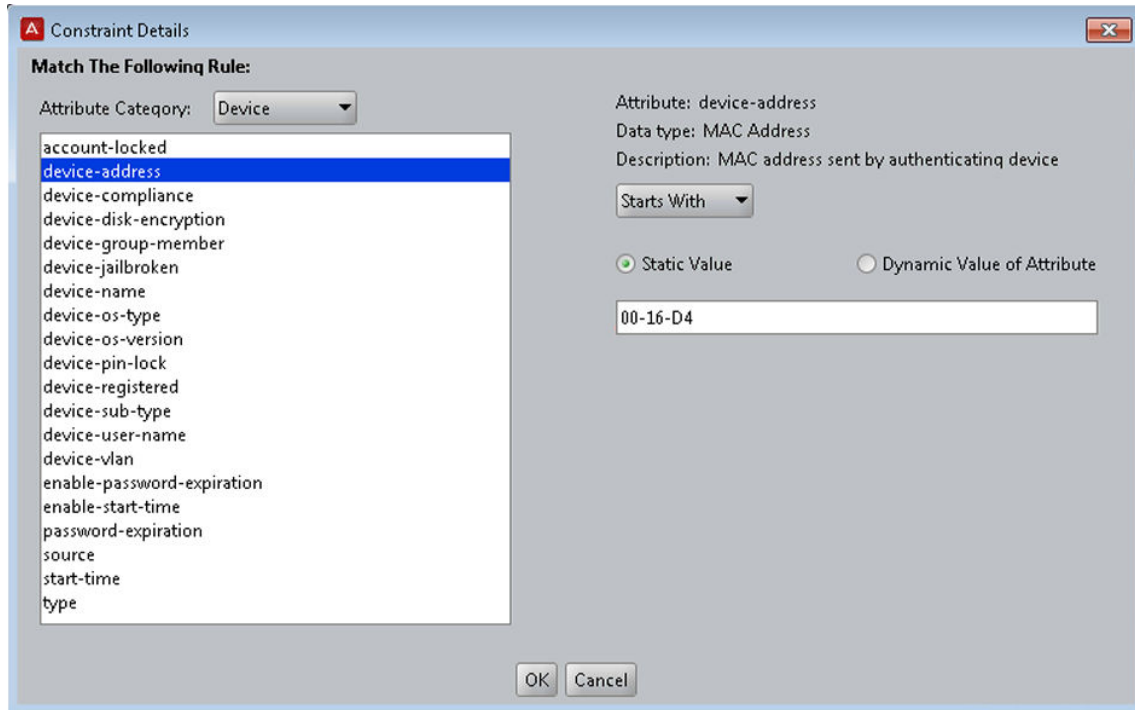
Procedure

1. In the **Configuration** tree, expand **Access Policies**.
2. Select **MAC Auth**.
3. Click **New**.

You can edit an existing policy by clicking its name in the **Configuration** tree and clicking **Edit**.

4. Enter a name for the policy and click **OK**.
5. Click the policy name in the tree and click **Edit**.
6. In the **Edit Authorization Policy** window, configure 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 264.

Typically, your MAC-Auth rules evaluate attributes of the connecting device. In the **Constraint Details** window, you configure 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.



For an example of a MAC authorization rule, see [MAC authentication set-up procedure example](#) on page 343.

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 234.

Setting Up MAC Auth

This section shows you how to enable MAC Auth for an authenticator. Later in this chapter, we provide an implementation example in the section, [MAC authentication set-up procedure example](#) on page 343.

Procedure

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 309.
2. Configure 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 configure these settings.
 - In the **RADIUS Settings** tab, select the **Enable MAC Authentication** check box.

- 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 340.)
- 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.

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 296.

- If you plan to perform VLAN assignment, select 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 135 or [Importing device records](#) on page 138.

 **Warning:**

Warning Concerning Users With MAC Addresses as Names

Under certain conditions, Ignition Server defies the precedence configured 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, MAC authentication should be enabled for the authenticator.

MAC authentication set-up procedure example

This example shows how to configure MAC authentication in Ignition Server. In this procedure, we build an policy example 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

1. Avaya recommends that, if you use MAC authentication, you configure a MAC Auth policy that assigns devices to one or more limited-access VLANs. To prepare for VLAN assignment.

- Configure 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 309. 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.



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 348.

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*.

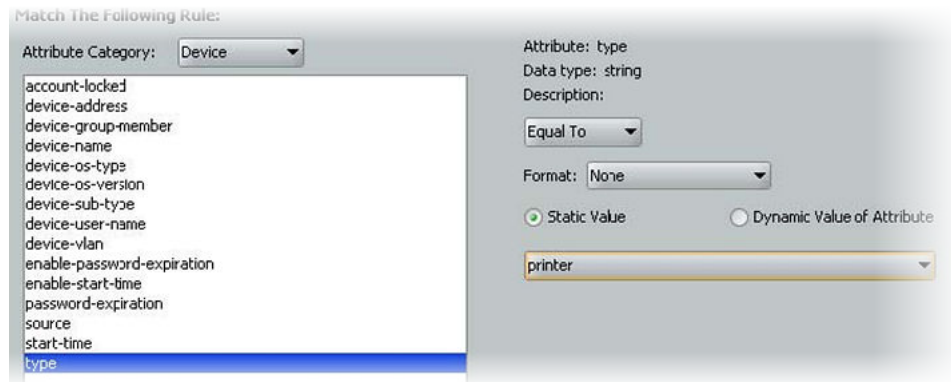
Ignition Server automatically checks that the device is a known device by checking the device's MAC address against the list of device records in the internal store.

The following steps show an example that performs VLAN assignment. To configure 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.

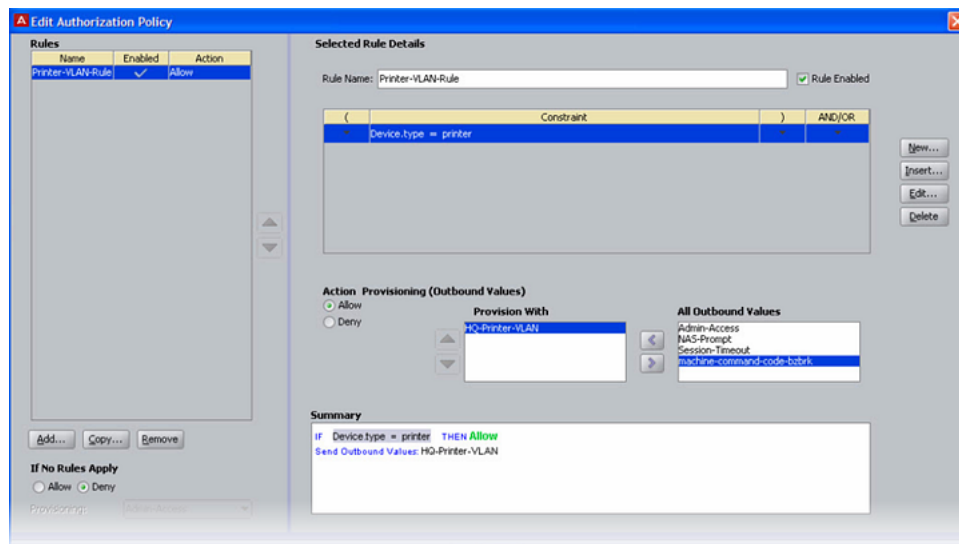
- In the **Edit Authorization Policy** window, click **Add** below the **Rules** list.
- In the **New Rule** dialog, give the rule a name and click **OK**. For example, you might call the rule, "Printer-VLAN-Rule",
- In the **MAC Authorization Policy** window, in the **Selected Rule details** section, 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** check box. In the drop-down list below this, click **printer** and click **OK**.

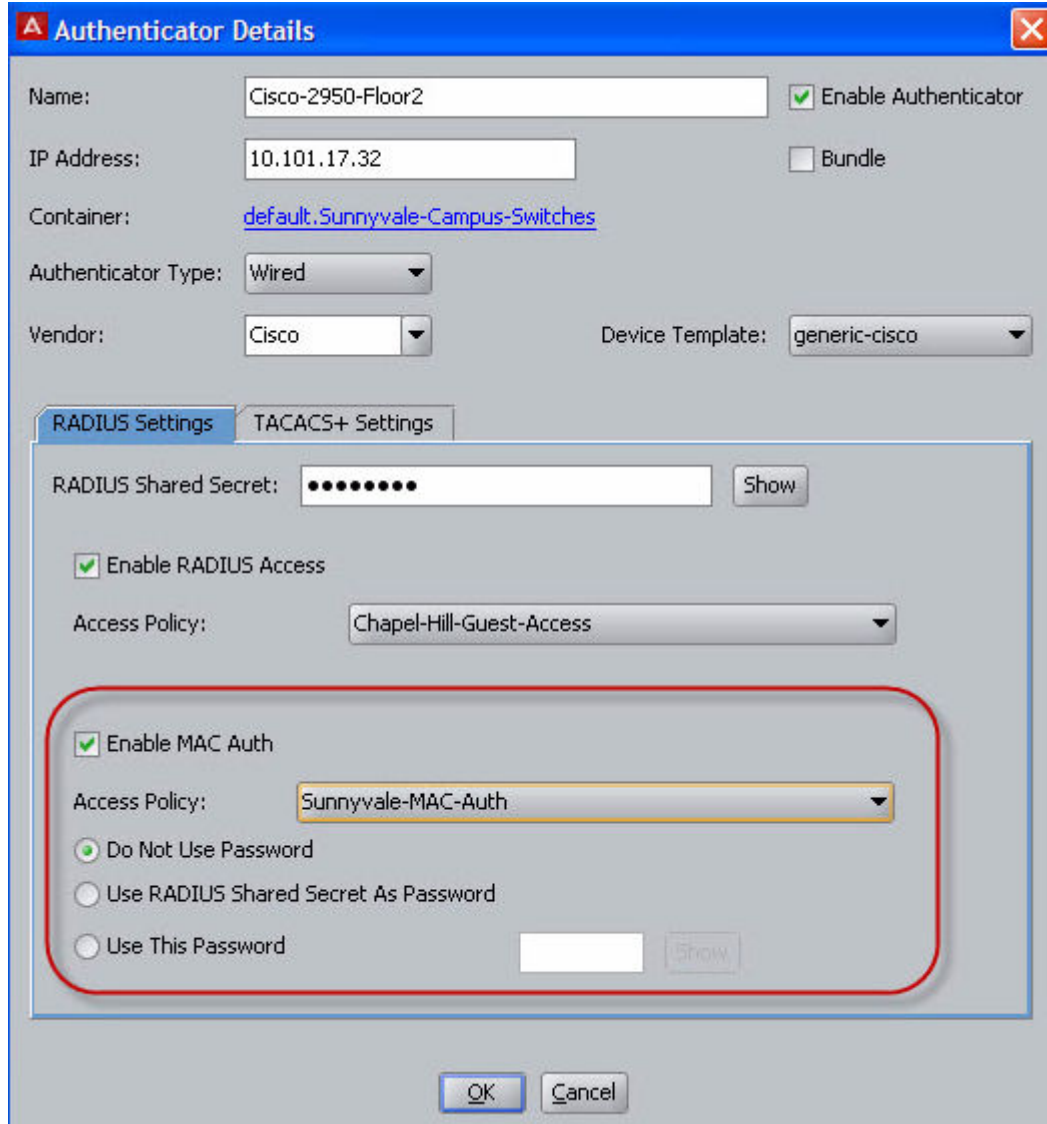
- In the **MAC Authorization Policy** window, with your "Printer-VLANRule" still selected, under **Action** select **Allow**. In the **Provisioning** section, select the check box next to "HQ-Printer-VLAN". (If this value is not in the list, create it now as explained at the beginning of this procedure.) Click **OK**.



- Your policy has been saved.

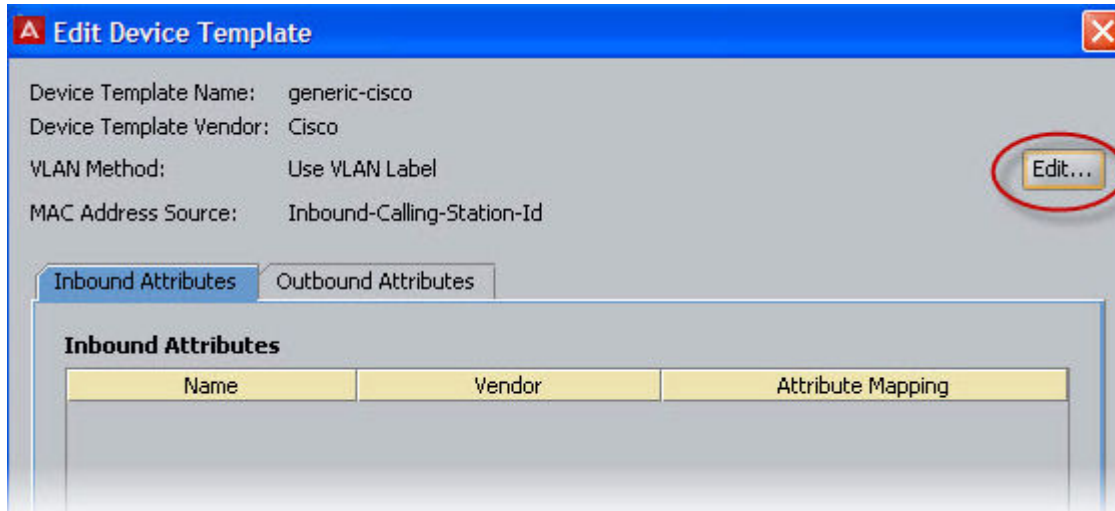


3. Configure the authenticators that support MAC authentication. Create or edit each authenticator record in the **Authenticator Details** panel of Dashboard. (From the main window, expand the Authenticators node in the hierarchy tree. Browse to find your authenticator, and click its name, and click **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 configured in Step 2.
 - Specify how the authenticator password should be checked. You have three choices. To skip password checking (the most common setting), select the check box **Do not use password**. To **use the authenticator's shared secret as the password**, select the check box **Use authenticator's shared secret as password**. To specify a password, select the check box, **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 296.
 - If you plan to perform VLAN assignment, select 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 135. For this example, make sure the **Type** of each device record is set to “printer”.

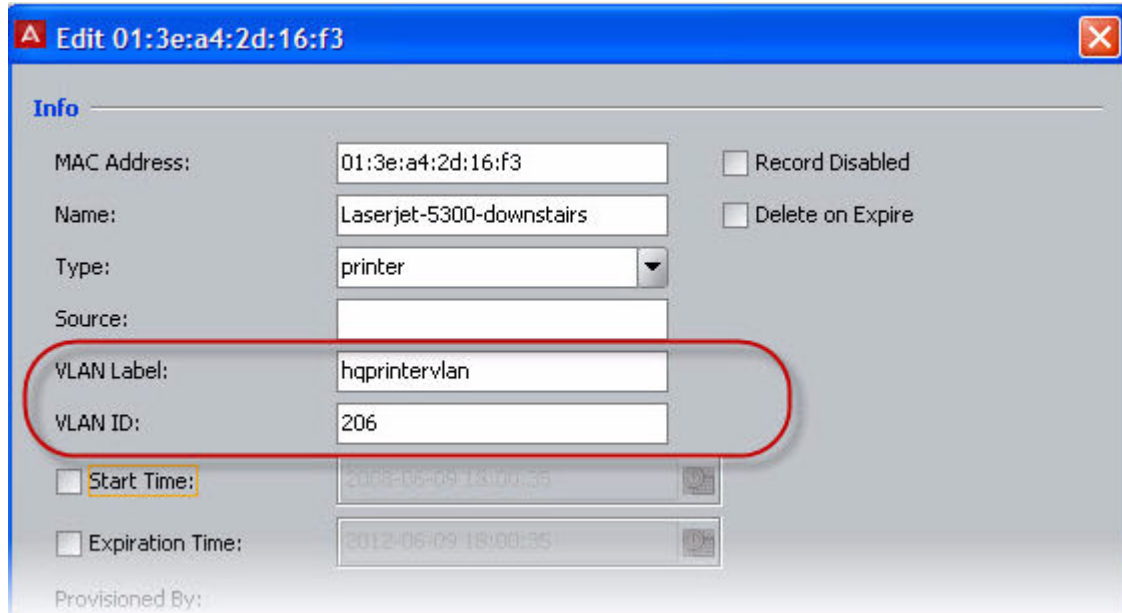
VLAN assignment using the Device Record VLAN fields

Using Ignition Server outbound value,s you can configure 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 [MAC authentication set-up procedure example](#) on page 343.)


Procedure

1. In the device templates of your authenticators, configure 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**.

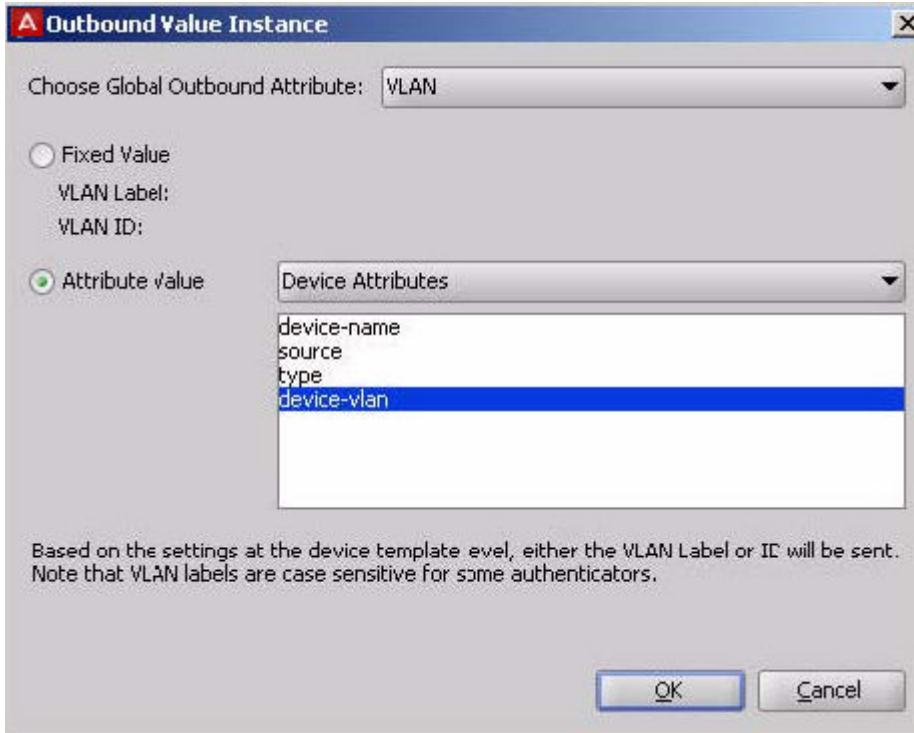
- In the **Edit Device Template** window, select the desired **VLAN Method**. **VLAN Label** uses a string and **VLAN ID** uses an integer value. Click **OK**.
2. In each device record, specify the desired VLAN.
 - In Dashboard's Configuration hierarchy tree, click your site, expand Site Configuration, 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.



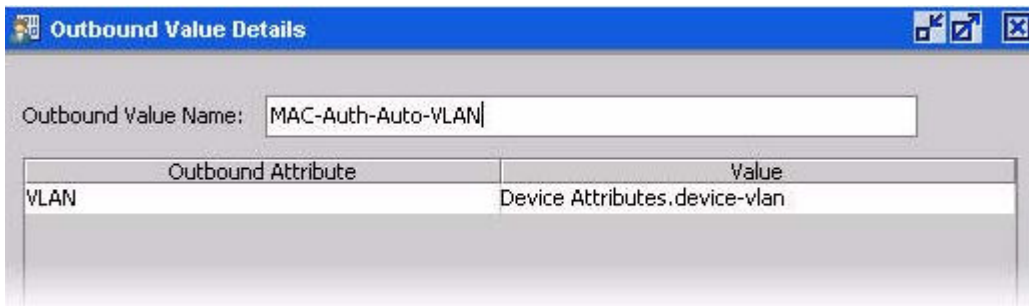
3. In Dashboard's Configuration tree, expand the Provisioning node and click **Outbound Values**.
4. In the **Outbound Values** panel, click **New**.
5. In the **Outbound Value Details** window, type a name for the outbound value. This is the name that appears in the **Constraint Details** window when 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**.
 - Select the **Attribute Value** check box 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 your 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.

The screenshot shows a configuration window with the following sections:

- Action:** Radio buttons for 'Allow' (selected) and 'Reject'.
- Provisioning (Outbound Values):** A list of checkboxes:
 - Admin-Access
 - NAS-Prompt
 - Session-Timeout
 - HQ-Printer-VLAN
 - MAC-Auth-Auto-VLAN (circled in red)
- Summary:** A text box containing the rule: `IF Device.type = printer THEN Allow With Outbound Values MAC-Auth-Auto-VLAN`
- Buttons:** 'OK' and 'Cancel' buttons at the bottom.

Allowed MAC Address formats

In the **Device Record** and **Constraint Details** windows, you can type a MAC address in any of the formats given in the following table. Upper and lower case letter characters are allowed. You can use colon, period, or hyphen characters as delimiters, but do not mix delimiters.

The following table lists the 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 286.
- **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

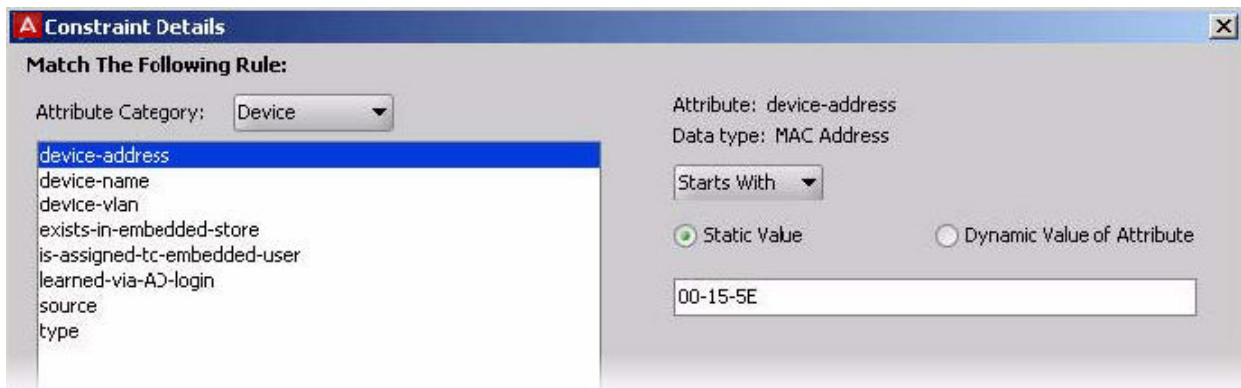
hierarchy, or the name of the Ignition Server MAC Auth access policy being used. See [Authenticator attributes](#) on page 260.

- **device attributes**: values that describe the connecting client device. See [Device Attributes](#) on page 259.

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**, select **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**), select **Static Value**, and click **Add** 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.

Chapter 22: 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.

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-embeddedstore** 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 selecting 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), delete its record from the Learned Devices tab of the Monitor: Current Site panel. To revoke a Windows-authenticated device’s right to connect, disable or delete its record in AD.

The following tables lists the 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.	
How to revoke the device’s right to connect?	Select the Device Disabled check box in its Ignition Server device record, <i>or</i> 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 355.

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 356.

OR

- that the device has authenticated itself to Ignition Server through *Windows machine authentication*, as explained in [Requiring the user to connect using a Machine Authenticated-Device](#) on page 358.

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.

Procedure

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**.

Configure your authentication policy as usual. ([User authentication policy](#) on page 237.)

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, as follows.

- 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-inembedded-store". On the right side of the window, click **False**. Click **OK**.

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.

To review the rule, 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 135 or [Importing device records](#) on page 138.

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 using the Device Record VLAN fields](#) on page 348.

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 234.

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.

If you use Windows machine authentication instead 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 358, instead of the following procedure.

Procedure

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. Configure your authentication policy as usual. ([User authentication policy](#) on page 237).
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 **False**. Click **OK**.

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.

To review your rules, 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 135 or [Importing device records](#) on page 138.
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 129.
8. Assign each user’s device to that user. See [Assigning a device to a user or group](#) on page 137.

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 using the Device Record VLAN fields](#) on page 348.

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 234.

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 when it performed its machine authentication.

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 an Allowed Device](#) on page 355, instead of the following procedure.

Follow this procedure to set up this asset correlation policy.

Procedure

1. Configure your machine authentication policy as explained in [Setting up Microsoft Windows Machine Authentication](#) on page 318.
2. Configure any VLANs you need.

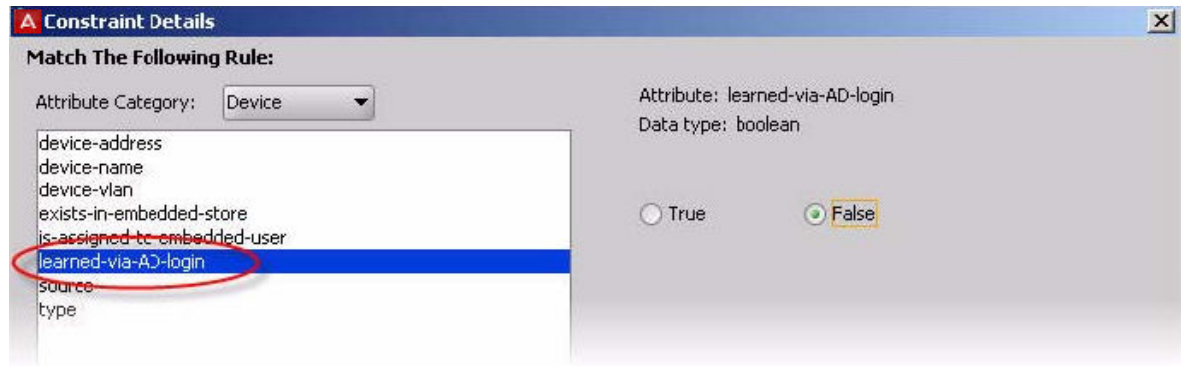
In this example we use two VLANs: LimitedAccess-VLAN which offers minimal access for users and machines that have not sufficiently authenticated, and HQ-Staff-VLAN which provides access to the internal network. To set up each VLAN:

- Configure 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 309.
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 **False**. Click **OK**.



- In the **Edit Authorization Policy** window, click the Action “**Allow**”, and under **Provisioning**, select the **Limited-Access-VLAN** check box and clear 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 “*Noprior-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**, select the **HQ-Staff-VLAN** check box and clear 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. Configure the TTL now as explained in [Setting TTL for Windows Machine authentication](#) on page 325. (In a running Ignition Server installation, you can view the current machine authentications as explained in [Viewing Currently Authenticated Devices](#) on page 360).

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 through *Windows machine authentication* appear in the **Learned Devices** tab. See [Learned Devices tab](#) on page 473.
- Devices authenticated through *MAC authentication* appear in the **AAA Summary** tab. See [AAA Summary tabs](#) on page 469.

Chapter 23: 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.

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 through 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.

If you have installed *no public keys* in the Ignition Server’s SSH service, the Server nonetheless allows you to connect through SSH. In this case, the Ignition Server authenticates you using only your system 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.

Procedure

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 password 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**.
 - Select the **Enabled** check box 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 administrator session over the local LAN with SSH encryption, do the following.

Procedure

1. Make sure that the following conditions are met.
 - 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 System 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 password. Enter your System 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.

Appendix A: Installing Ignition Server

This appendix explains how to install Ignition Dashboard and how to connect and configure the Avaya Identity Engines Ignition Server.

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.

- 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 System administrator name ("admin") and password ("admin")
- an IP address and subnet mask you can assign to each Ignition Server network port. See [Configuring the Ignition Server's network ports](#) on page 71. 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 require the types of information listed as follows.

*** Note:**

If you are performing a basic installation, you need not gather this information now. Instead, follow the steps in [InstallingTheIgnitionDashboardDesktopApplication](#) on page 378.

- 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
- 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 5.1 and up are supported. The VM requires an x86_64 capable environment, a minimum of 4 GB of memory, a minimum of 250 GB of available disk storage (thin provisioning is allowed), a minimum of four 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 OVA file, which already incorporates the OS Red Hat Enterprise Linux.

Reminder: Avaya provides the Identity Engines Ignition Server, Ignition Guest Manager, 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

Use this procedure to prevent automatic VMware Tools updates.

Procedure

1. Use the VMware vSphere Client to log in to the ESXi Server hosting the Ignition VM.
2. Select the VM corresponding to the Ignition Server.
3. 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.
4. Click **OK**.



Checking the VMware Tools status (ESXi 5.1 and up)


The **Summary** tab of the VM describes the VMware Tools status.

To check the VMware Tools status on an ESXi (5.1 and up) server:

Procedure

1. Use the vSphere client to log in to the ESXi Server.
2. Go to the **Summary** tab.

After a fresh install, the VMware Tools status displays as “VMware Tools: Running (Current)”.



*** Note:**

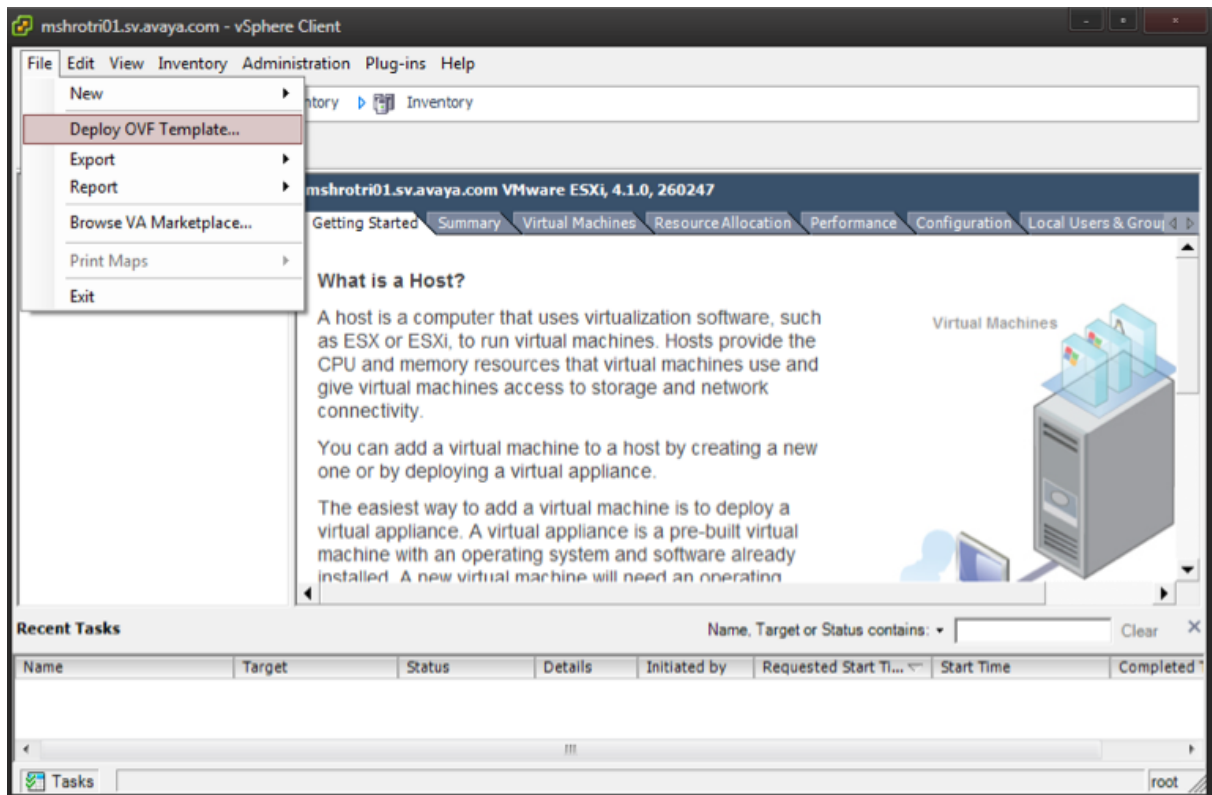
VMware Tools may show as not installed. This is a known VMware issue where VMware Tools may not be detected correctly on certain hardware. However, this does not interfere with the functioning of the tools—it is a display issue only.

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.

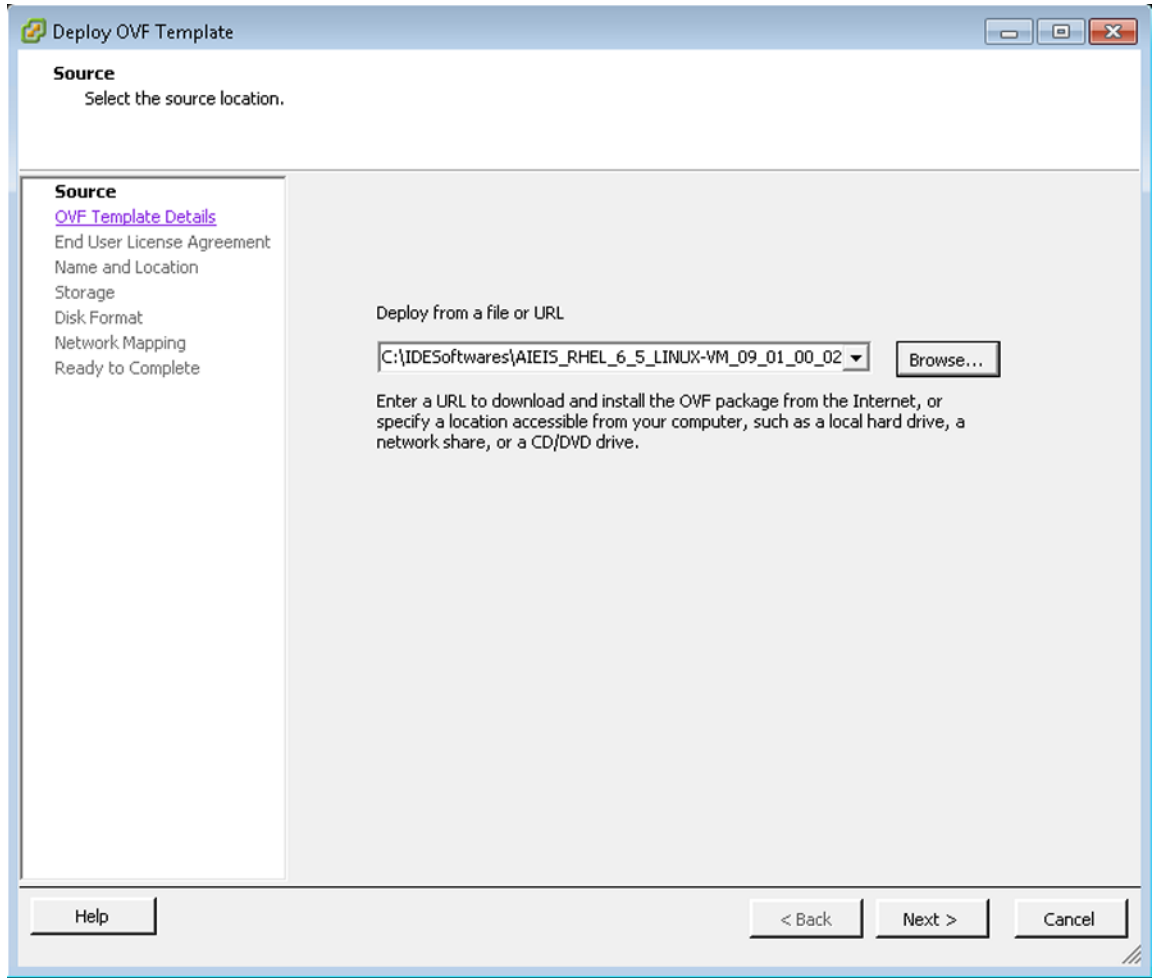
Procedure

1. From the vSphere Client, select **File > Deploy OVF Template**.



The **Source** screen displays.

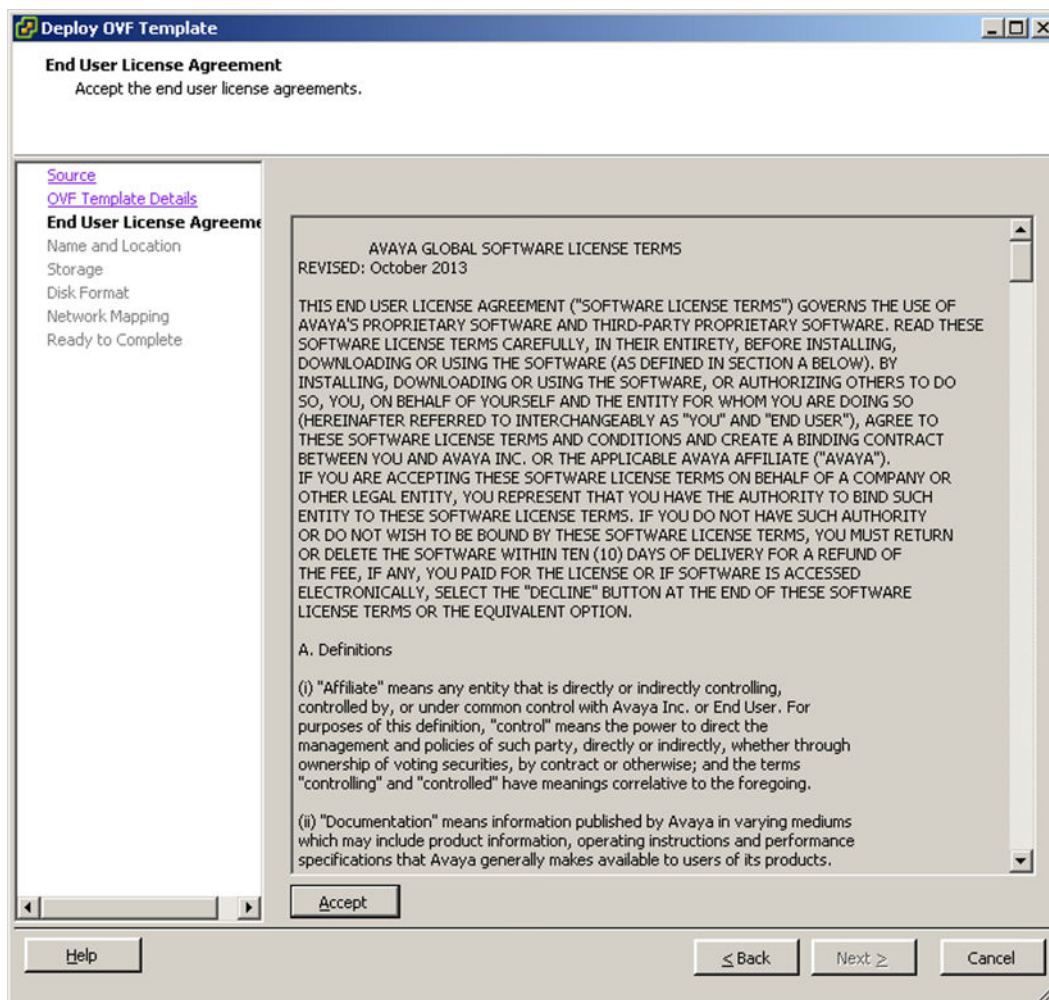
2. Select the location from which you want to import the Ignition Server virtual appliance.



3. Click **Next**. In the **OVF Template Details** screen, review your settings. Click **Back** to make changes, or click **Next** to continue.

The **End User License Agreement** screen displays.

4. Click **Accept** to accept the licence and click **Next**.




The **Name and Location** screen displays.

5. Either accept the default name or choose to rename the virtual machine. Click **Next**.

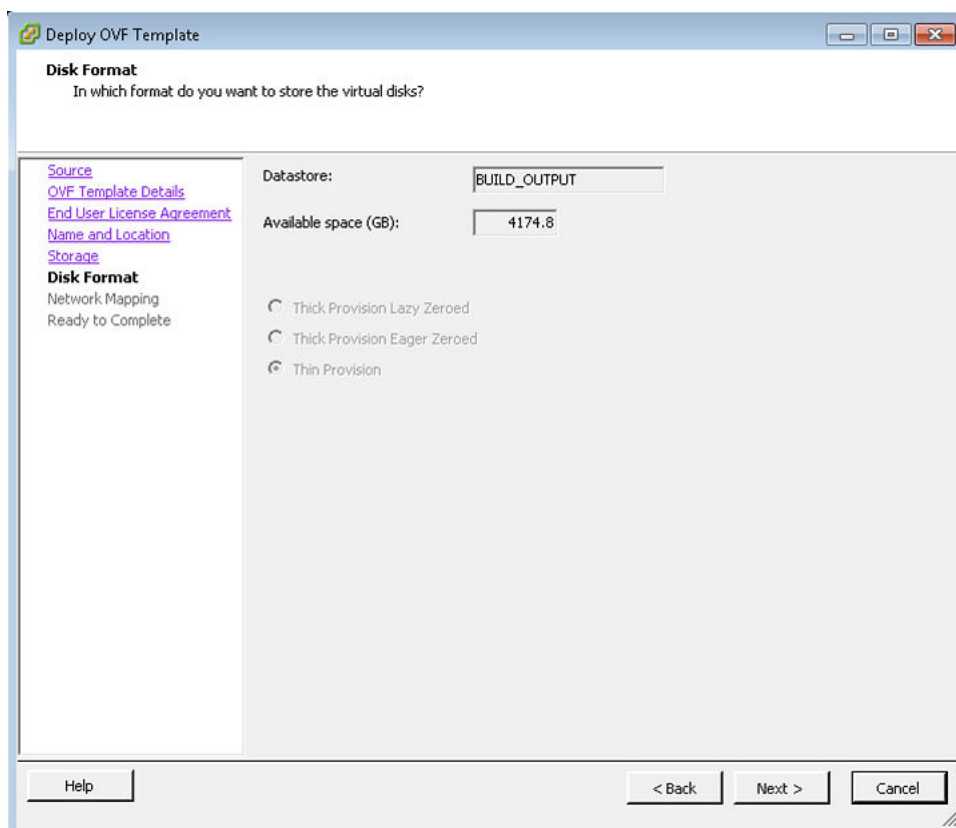
The **Datastore** screen displays.

6. Select the location where you want to store the files for the virtual appliance and click **Next**.



The **Disk Format** screen displays.

7. Select a format in which to store the virtual machine's virtual disks and click **Next**.



The **Network Mapping** screen displays.

8. Associate the Avaya Ignition Server NIC's to the correct VM Network based on your site configuration. Then click **Next**.

The **Ready to Complete** screen displays.

9. Review your settings. Click **Back** 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 must verify and adjust some of the VM settings. Open the **VM setting** dialog and select the **Options** tab. Do the following:
 - a. 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 correct NIC for each interface.

You are now ready to boot the Avaya Ignition Server for the first time. A splash screen displays as the boot up starts.

11. Once the Ignition Server Console login prompt displays, you are ready to enter the administration IP address. Log in using `admin` for the user name and `admin` for the password. Avaya recommends that you should change the password after you login.

```

Applying Intel CPU microcode update: [ OK ]
Performing Avaya Ignition Server Network Setup
Starting background readahead: [ OK ]
Checking for hardware changes [ OK ]
Starting system logger: [ OK ]
Starting kernel logger: [ OK ]
Starting irqbalance: [ OK ]
Starting vmware-tools: Starting VMware Tools services in the virtual machine:
  Switching to guest configuration: [ OK ]
  Paravirtual SCSI module: [ OK ]
  Guest memory manager: [ OK ]
  VM communication interface:UMCI: Major device number is: 253
  [ OK ]
  VM communication interface socket family: [ OK ]
  Guest operating system daemon: [ OK ]
  [ OK ]
Starting system message bus: [ OK ]
Starting xinetd: [ OK ]
Starting xfs: [ OK ]
mount: block device /dev/loop0 is write-protected, mounting read-only
Starting Avaya Ignition Server:

Ignition Server Console
login: _

```

12. Use the interface commands as shown in the following 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> “

```

Starting irqbalance: [ OK ]
Starting vmware-tools: Starting VMware Tools services in the virtual machine:
  Switching to guest configuration: [ OK ]
  Paravirtual SCSI module: [ OK ]
  Guest memory manager: [ OK ]
  VM communication interface:UMCI: Major device number is: 253
  [ OK ]
  VM communication interface socket family: [ OK ]
  Guest operating system daemon: [ OK ]
  [ OK ]
Starting system message bus: [ OK ]
Starting xinetd: [ OK ]
Starting xfs: [ OK ]
mount: block device /dev/loop0 is write-protected, mounting read-only
Starting Avaya Ignition Server:

Ignition Server Console

login: admin
password:
Ignition Server> interface admin ipaddr 134.177.229.200/24
Success: Interface admin's ipaddr/netmask is set to 134.177.229.200/24.
Ignition Server> interface admin enable
Success: Interface admin is Enabled.
Ignition Server> _

```

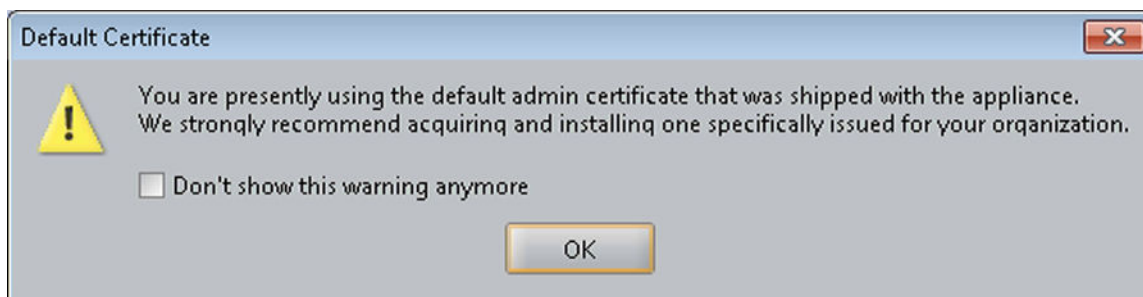
13. Install the Dashboard on to your Desktop machine.

See [InstallingTheIgnitionDashboardDesktopApplication](#) on page 378.

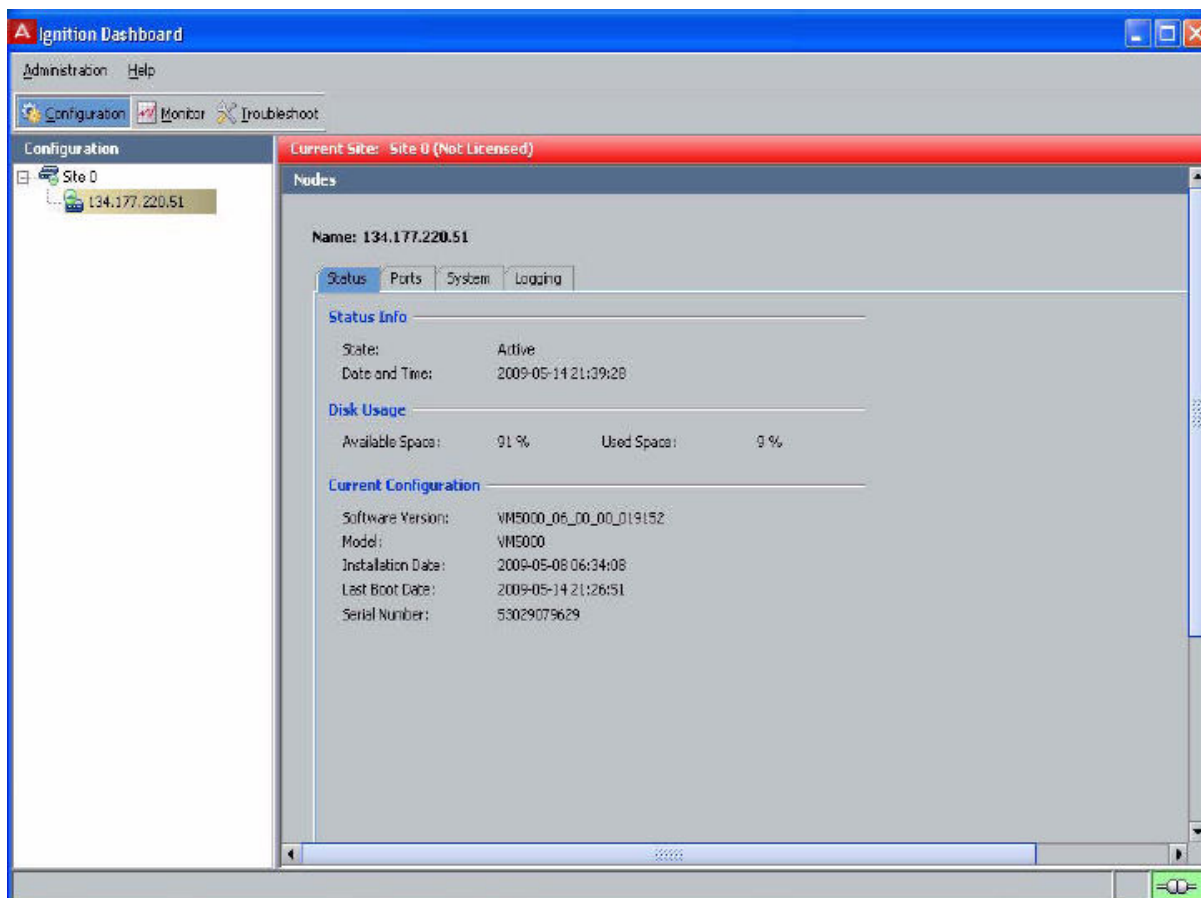
14. Once installation is complete, click on the desktop icon to start the application.

A **Login** dialog displays.

15. Enter 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 see the following message.



If you click **OK** to both dialogs, you see a display similar to the following figure.



Next steps

In order to obtain your license, see [Applying the license](#) on page 376. Once you have obtained your license, you can proceed with the final configuration of the Avaya Ignition Server in your environment.

Applying the license

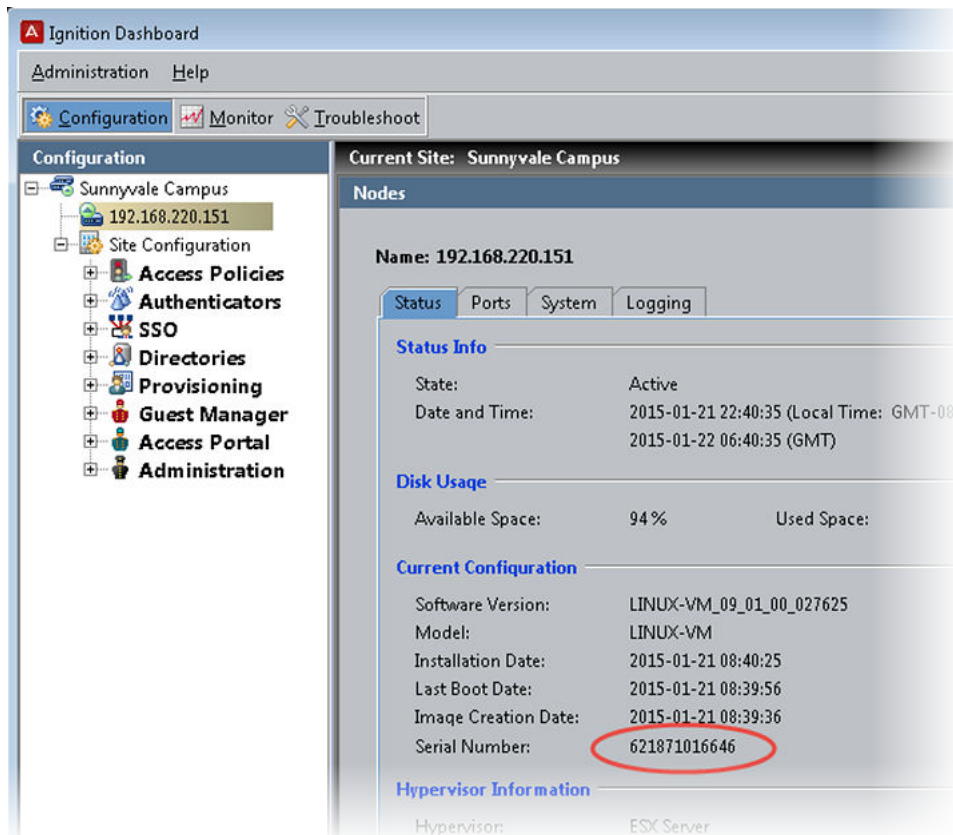
The Avaya Identity Engines Ignition Server (AIEIS) Software ships without any licenses. There are seven different software licenses that can be installed on Ignition Server: Base License, Guest Manager License, NAP Posture License, TACACS+ License, Ignition Reports License, Access Portal License, and Avaya Aura® Single-Sign-On (SSO) 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).

Procedure

1. 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/>.
2. Once this is purchased, Customer Support sends a software CD and certificate that contains a unique product code and an email address. Send this unique product code and the Node Serial Number to the email address provided. The Node Serial Number can be obtained from Dashboard from the Status tab of Node Configuration as shown in the following figure.



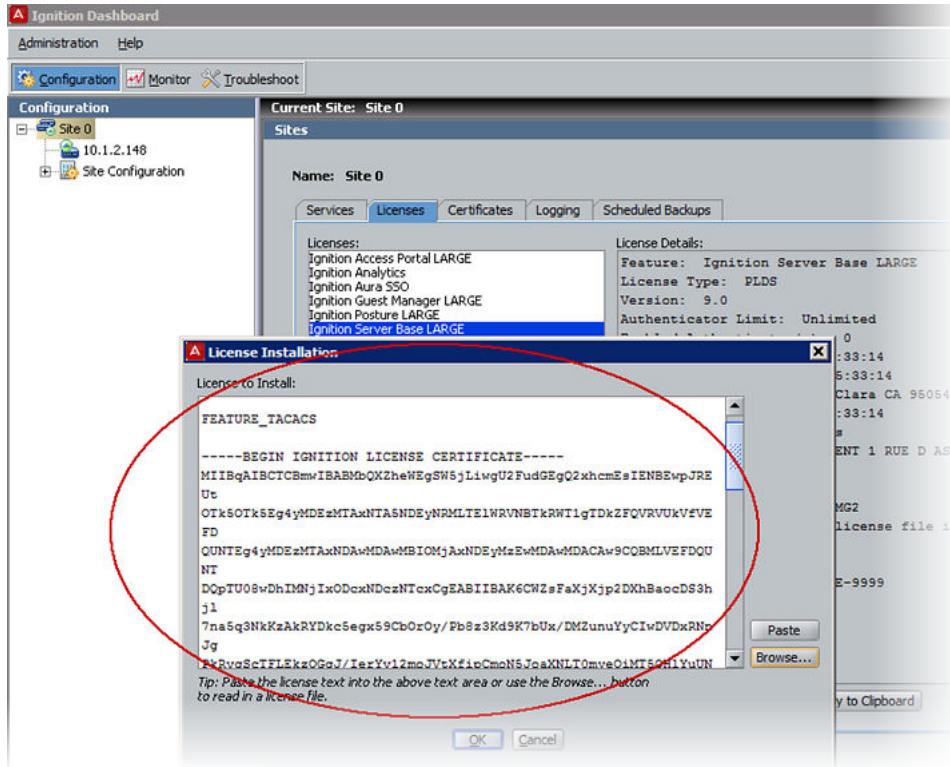
3. After the unique product code and Node Serial Number is verified, a software license file is sent back to you. Install this license on the server using Dashboard.

Installing the license

You can install the license on the Ignition Server using Dashboard. To install the license, perform the following steps:

Procedure

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**.



Installing the Ignition Dashboard desktop application

The Avaya Ignition Dashboard is a desktop application that enables you to manage the Ignition Server appliance. The Avaya Ignition Dashboard enables you to create, view, or alter configuration information for authenticators, service categories, and the policies that apply to authentication and authorization.

Before you begin

To proceed with the Ignition Dashboard installation, have the following tools and information ready:

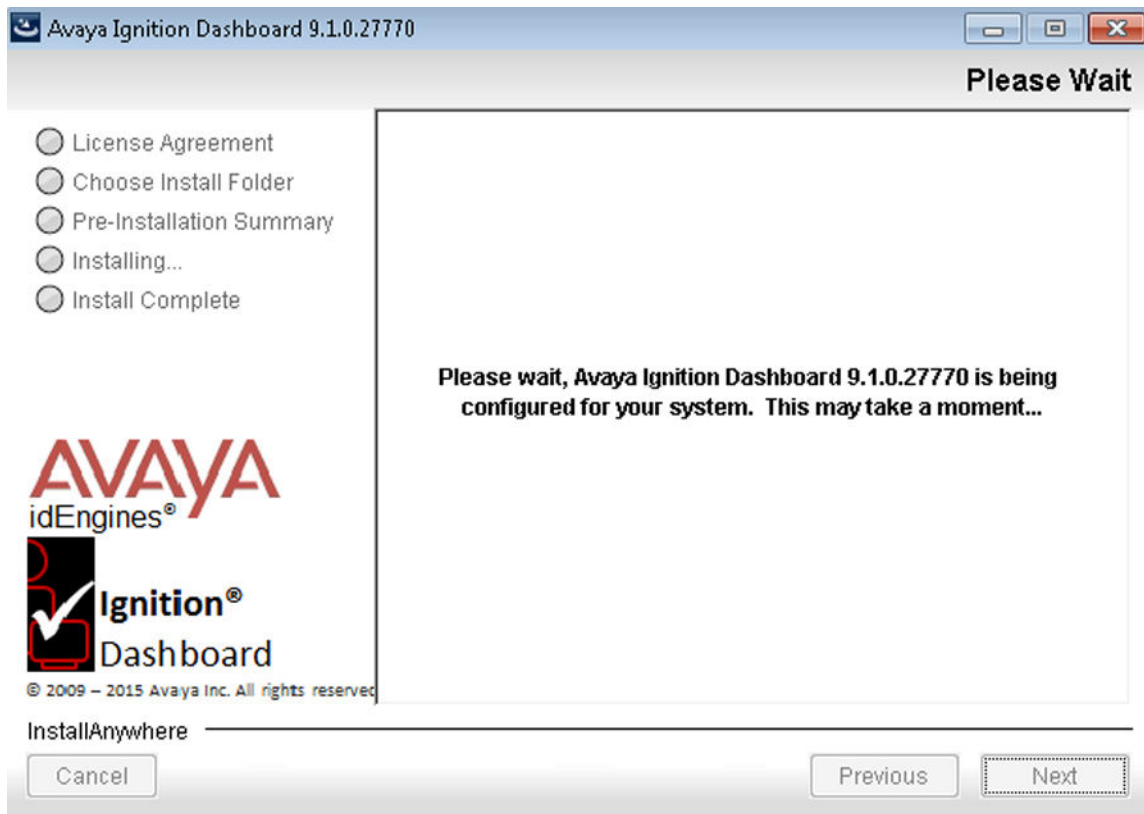

- The Identity Engines product software shipped with your Ignition Server appliance.
- A computer running Windows 7 (32 bit or 64 bit), Windows 8 (32 bit or 64 bit), or Windows Server 2008 (32 bit or 64 bit).
- A minimum of 2 GB of RAM memory.
- The default System administrator name (`admin`) and password (`admin`).

Procedure


1. If any version of the Avaya Ignition Dashboard exists on the computer, ensure the Ignition Dashboard application is not currently running. If the Ignition Dashboard is running, shut it down now.

- Place the Ignition Server CD into the CD drive of your computer. On Windows, the Windows AutoRun feature runs the Installer immediately.

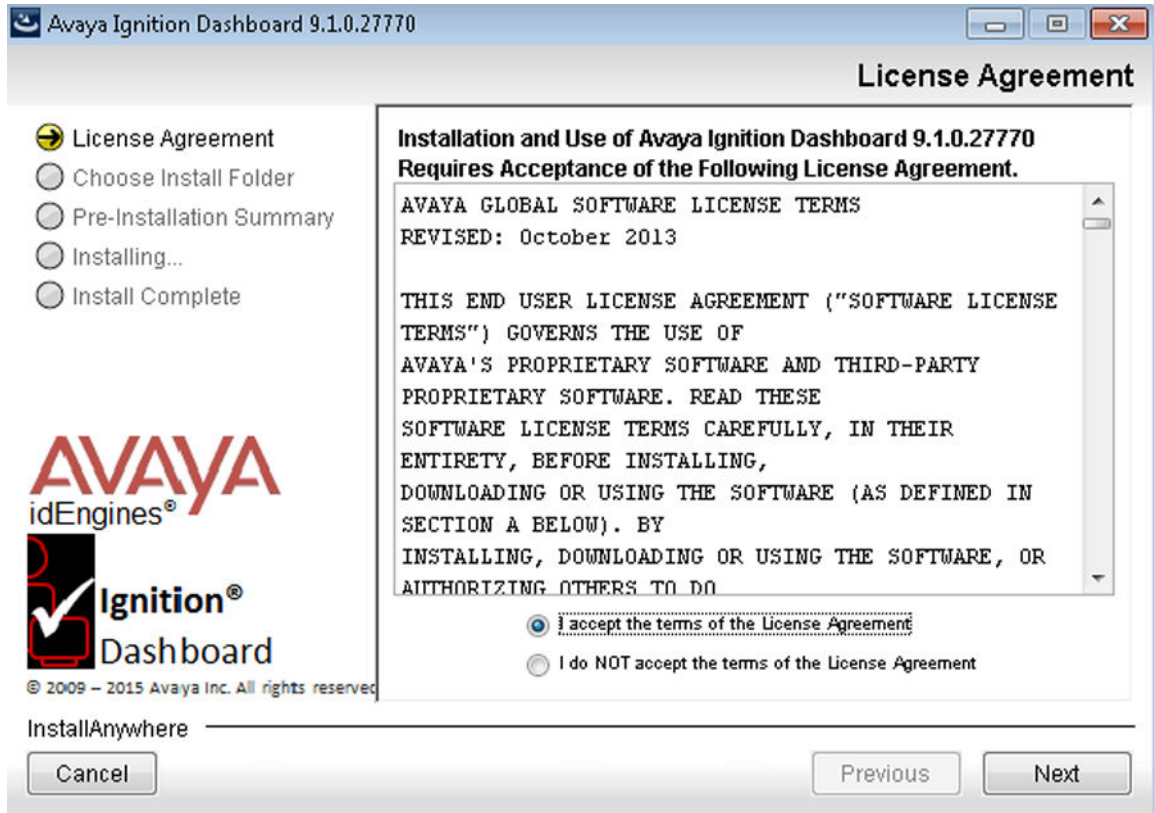
Note: If the AutoRun feature is disabled on your computer, navigate to your CD drive and double-click the installer file. It has a name like DashboardInstaller-9.1.0<Build Number>.exe.




- If an older version of Ignition Dashboard exists on your device, the **Existing Dashboard Found!** window displays. To remove the old Ignition Dashboard, select **Yes**. To install the new version of Ignition Dashboard without removing the older version, select **No**.



4. In the **License Agreement** screen, scroll down to read the entire license. Select the radio button to accept the license and click **Next**.



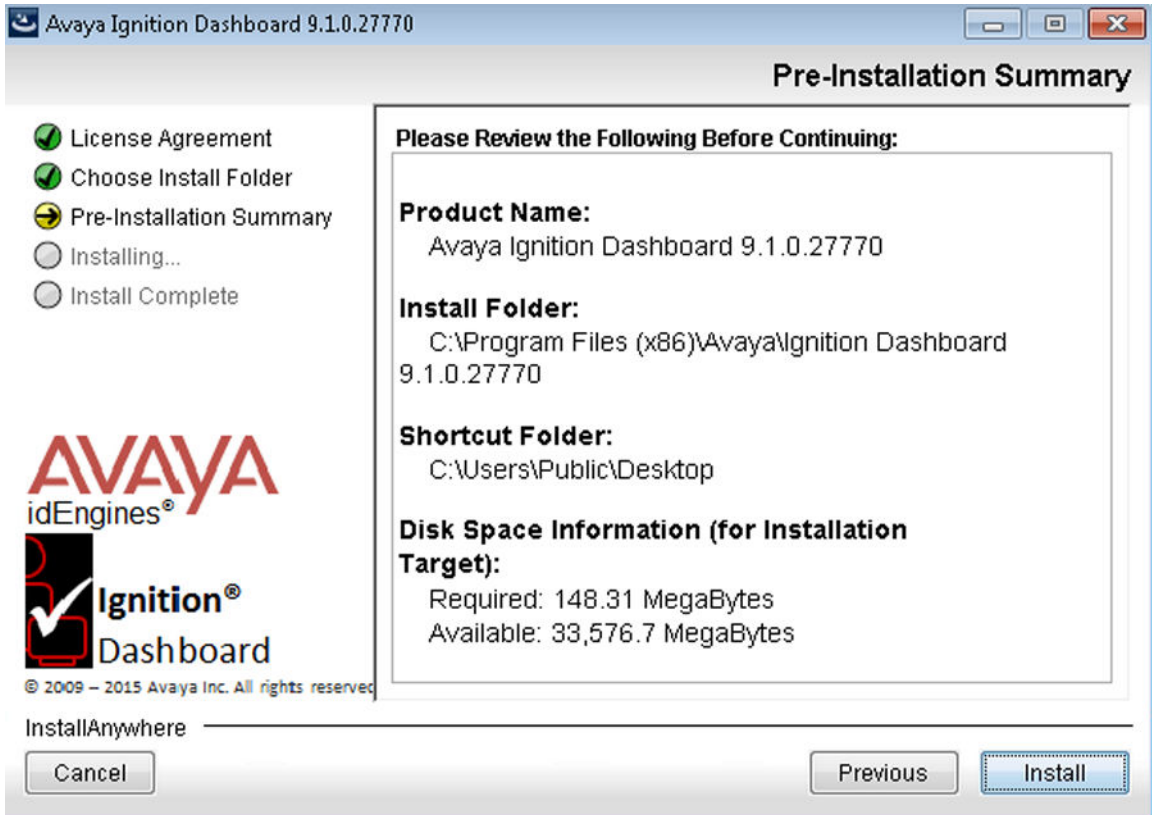
5. In the **Choose Install Folder** screen, choose your destination folder and click **Next**.



6. In the **Choose Shortcut Folder** screen, indicate where you want the Dashboard shortcut to appear, and click **Next**.



7. In the **Pre-Installation Summary** screen, review your installation settings. If you want to make changes, click **Previous** to edit the details of the locations of the installation. When you finish your configuration, click **Install**. The installer displays a pre-install confirmation window.




8. In the **Pre-Installation Summary** confirmation window, click **OK** to confirm the installation.


! Important:

Ignition Dashboard requires JRE 1.7.0_45 to operate correctly. For this reason, during the Dashboard installation, JRE 1.7.0_45 is installed in static mode. For information on JRE static installation, see the “Static Configuration” section of the JRE installation at <http://docs.oracle.com/javase/7/docs/webnotes/install/windows/jre-installer-options.html>.

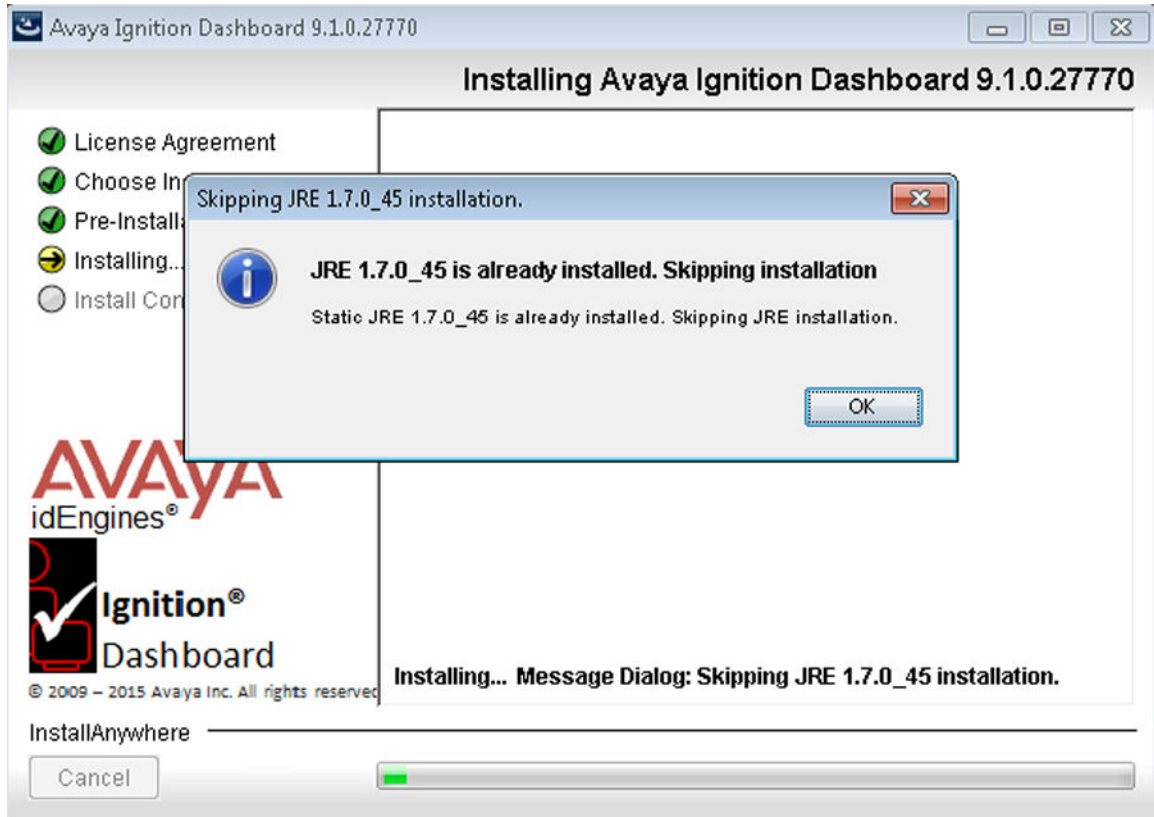
You can upgrade the system JRE that is installed on your machine by default, but do not delete the JRE 1.7.0_45 version from your machine either manually or as part of the JRE upgrade. If you delete JRE 1.7.0_45, the Ignition Dashboard application will not function as expected and results in unpredictable behavior.



The installation starts. The installer displays a dialog box that displays the progress of the installation.

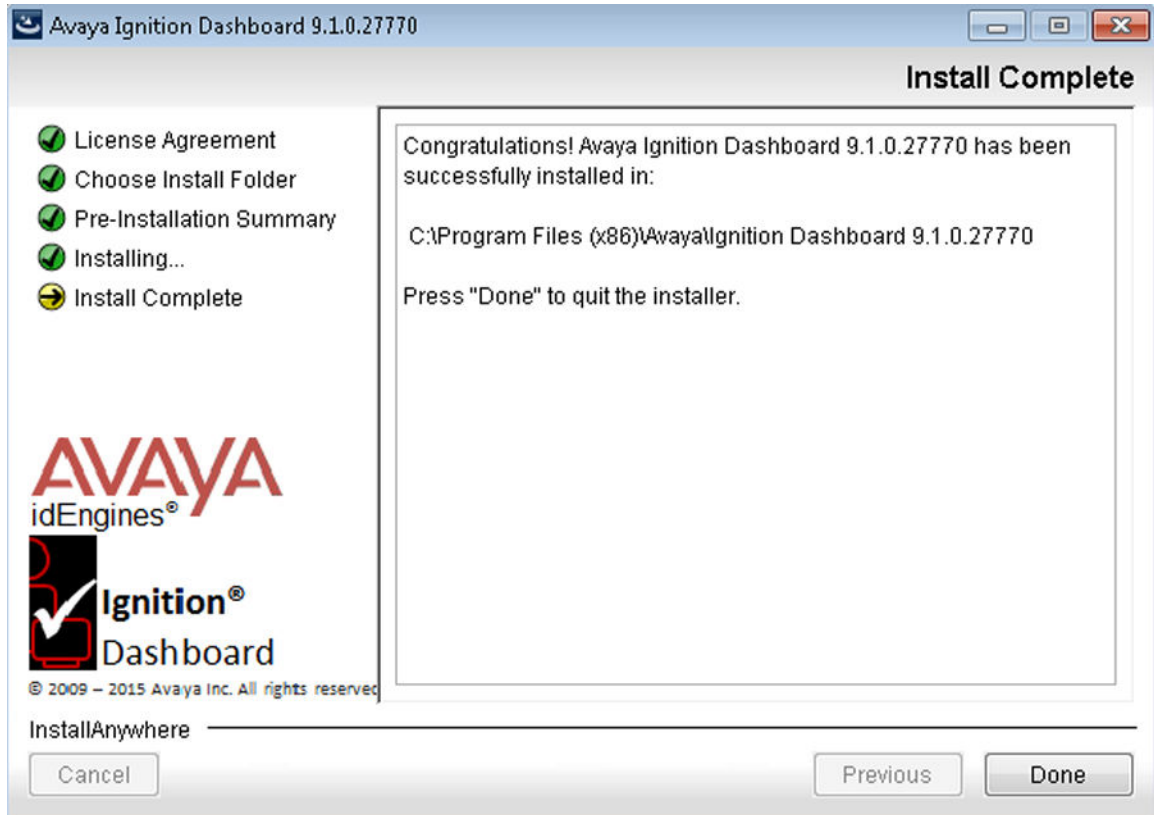


9. If the appropriate Java Runtime Environment (JRE) is already installed, a window appears to allow you to skip re-installing JRE. Click **OK** to skip the re-installation of JRE.



The installation continues. The installer displays a dialog box showing the progress of the installation.

10. When the installation is complete, the installer displays the **Install Complete** screen. In the **Install Complete** screen, click **Done**. An icon for Ignition Dashboard appears in the location you designated.



Installing multiple versions of the Ignition Dashboard: You can 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.

Connect Ignition Server for the first time

Run Ignition Dashboard

Procedure

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 479.
- If your login attempt succeeds, a warning dialog displays 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 88.

After you dismiss the warning dialog, Ignition Dashboard appears.

Change the System Administrator Password

Avaya recommends you change the default password when you first set up your Ignition Server. To change the System administrator password, see [Configuring the System Administrator password](#) on page 57.

Next steps for installers: Your Ignition Server installation is complete. See [Configure the Ignition Server](#) on page 389 for a list of configuration options.

Install Your Ignition Server Licenses

See [Installing an Ignition Server license](#) on page 78.

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.

Object	Reference Chapter
Authenticators: Representing wireless access points, Chapter , groupings of wireless access points, or network devices “Authenticators” such as VPN concentrator or server, Ethernet switches, WLAN switches, or routers.	Authenticators on page 97
Directory services: Repositories of user identities and Chapter , “Directory attributes such as Active Directory, LDAP, and token Services” servers. Directory sets: Groups of directory services.	Directory Services on page 148
Users, groups, and attributes: Entities or objects Chapter , “Internal represented in directories or databases, that contain Users, Groups, and information about end users. Devices” .	Internal users, groups, and devices on page 127

Table continues...

Object	Reference Chapter
<p>Authentication and authorization policies: Each access Chapter , “User policy establishes a set of rules that governs user access. Authentication Policy” Rules are evaluated based on user attributes and other criteria.</p>	<p>User authentication policy on page 237</p>
<p>Provisioning policies: Optionally, each access policy Chapter , may have a provisioning policy that assigns each user to “Provisioning Policy” an appropriate VLAN and/or sets switch parameters for the user.</p>	<p>Provisioning policy on page 274</p>

Uninstalling Ignition Dashboard

Follow this procedure to uninstall Ignition Dashboard.

Procedure

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.

Appendix B: Paired server 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 and/or SAML 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.

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 and/or SAML 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 the *database primary node* serves configuration requests (it handles the data changes the administrator submits). The Ignition Server acting as the *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, see [Warning for Sites Running Ignition Server in HA Mode](#) on page 204.

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 they 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 [Setting up logging](#) on page 438 and [Viewing logs and statistics](#) on page 450 for instructions.

Network settings: The System Administrator configures 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 configuration.

Creating an HA Pair

Use the following procedure to create your HA pair and enable the VIP that provides failover for Ignition Server services.

Start and connect the Ignition Server

Procedure

1. Configure and start both Ignition Server nodes, as explained in the *Avaya Identity Engines Ignition Server Getting Started Configuration, NN47280–300*.

Avaya strongly recommends that you use two virtual machines on two different servers to ensure that the HA can maximize 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 Layer 2 switch, as this provides support for high-availability environments. Make sure that your port network connections comply with the following rules.

- 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.
 - The port subnet must be reachable from your authenticators and from your Ignition Dashboard workstation.
 - The network of these connections 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 ERS switches or the Avaya VSP switches is appropriate.
3. On the first node, use the Ignition Server dashboard to configure network settings for the Admin port. Configure the IP address, subnet mask and gateway address. For help using the Ignition dashboard, see *Avaya Identity Engines Ignition Server Getting Started Configuration, NN47280–300*.
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 configure 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 IP 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, use the following procedure.

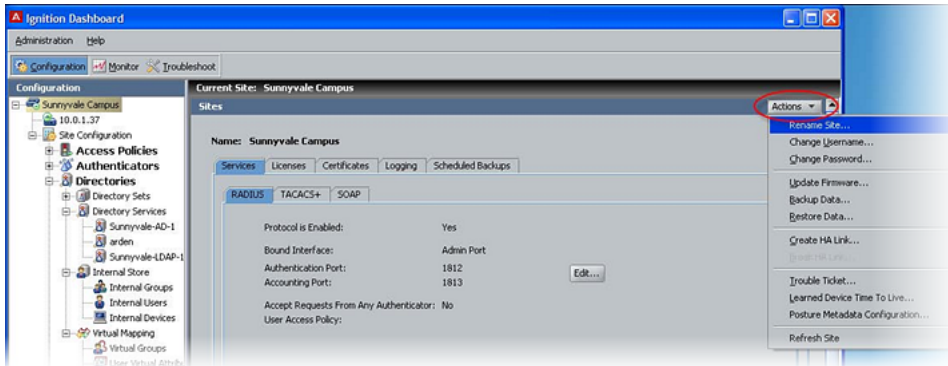
Before you begin

- **Secondary node data will be erased:** Designate one node as primary and one node as secondary for the duration of the configuration 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.

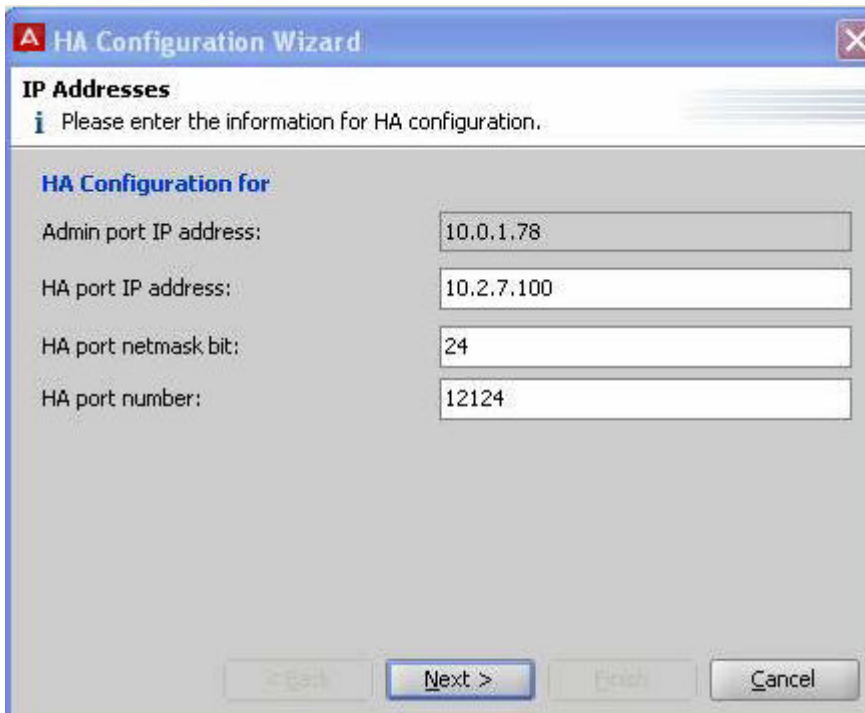
This procedure continues from Step 6 in the preceding procedure.

Procedure

1. Using Ignition Dashboard, log in to either of the Ignition Servers 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 displays.



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 IP 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** displays, requiring login information for the second node in the HA pair (the second Ignition Server in the pair).


The screenshot shows a window titled "HA Configuration Wizard" with a "Login" section. Below the title bar, there is an information icon and the text "Input the secondary HA unit's Access Login Information below." The form contains three input fields: "Username" with the value "admin", "Password" with five dots, and "Hostname" with the value "10.0.1.78". At the bottom of the window, there are four buttons: "< Back", "Next >", "Finish", and "Cancel".

Enter the System 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**.

- The HA Configuration Wizard requires you designate the primary node for the newly created HA pair. Select the required primary node in the new HA pair, and click **Next**.
- If prompted, you must specify which node's System 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 configured to match those of the node you select. This is the only Dashboard login for the pair of Ignition Servers. You can change the administrator password later, as explained in [Configuring the System Administrator password](#) on page 57.

- The HA Configuration Wizard asks whether you want to create a VIP. Select **Yes** and click **Next**.



- Configure the VIP settings in the **Virtual Interface Configuration** window.

The settings are explained as follows, after the illustration. (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 **Delete all existing virtual interface definitions...** and click **Next**. If you want to restore the VIP, see [Restoring a saved VIP configuration](#) on page 402.)

HA Configuration Wizard

Virtual Interface Configuration

i Create a virtual interface and bind it to RADIUS.
The VIP should be used when connecting to your HA pair.

Name:

Virtual Host ID: **i**

Password: **i**

VIP IP Address: /

Bind To:

Enabled:

< Back Next > End Cancel

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 Servers in the pair.

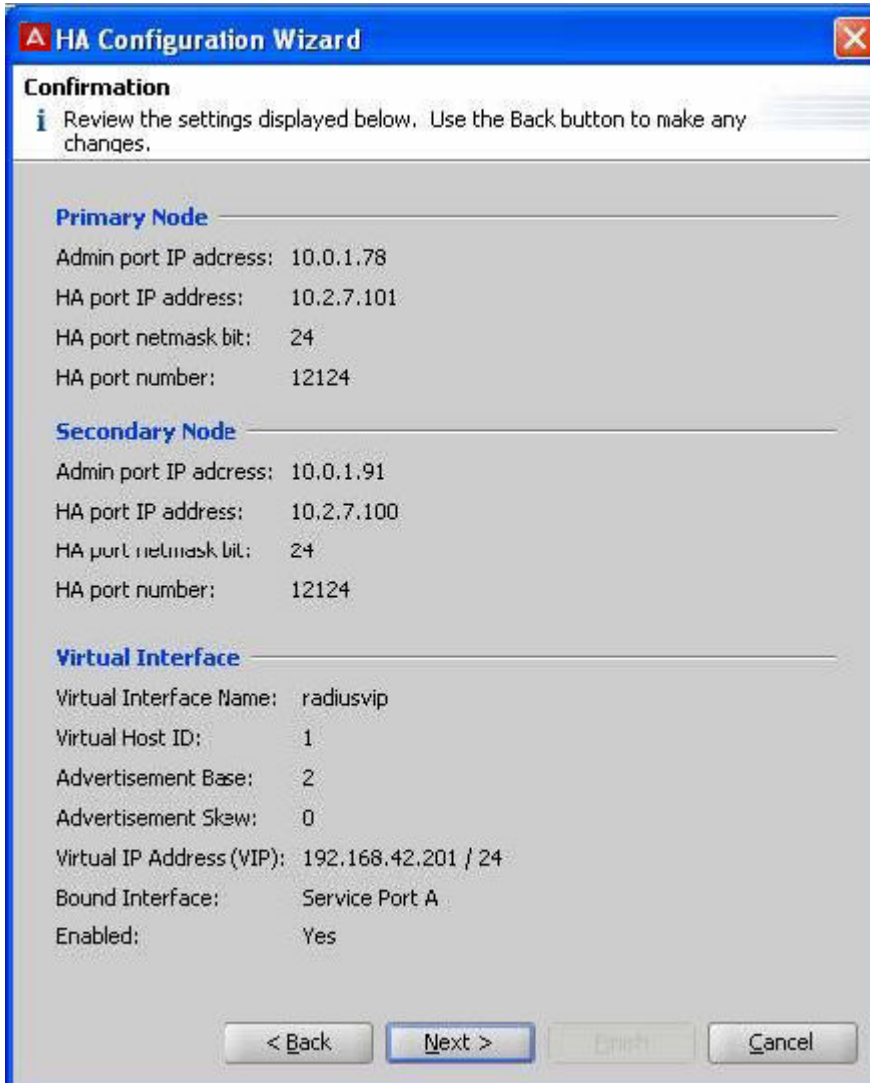
! Important:

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:** Select this checkbox to enable the virtual interface.

For complete field descriptions, see [Step 4](#) on page 406. For general VIP information, see [Managing Virtual Interfaces \(VIPs\)](#) on page 405.

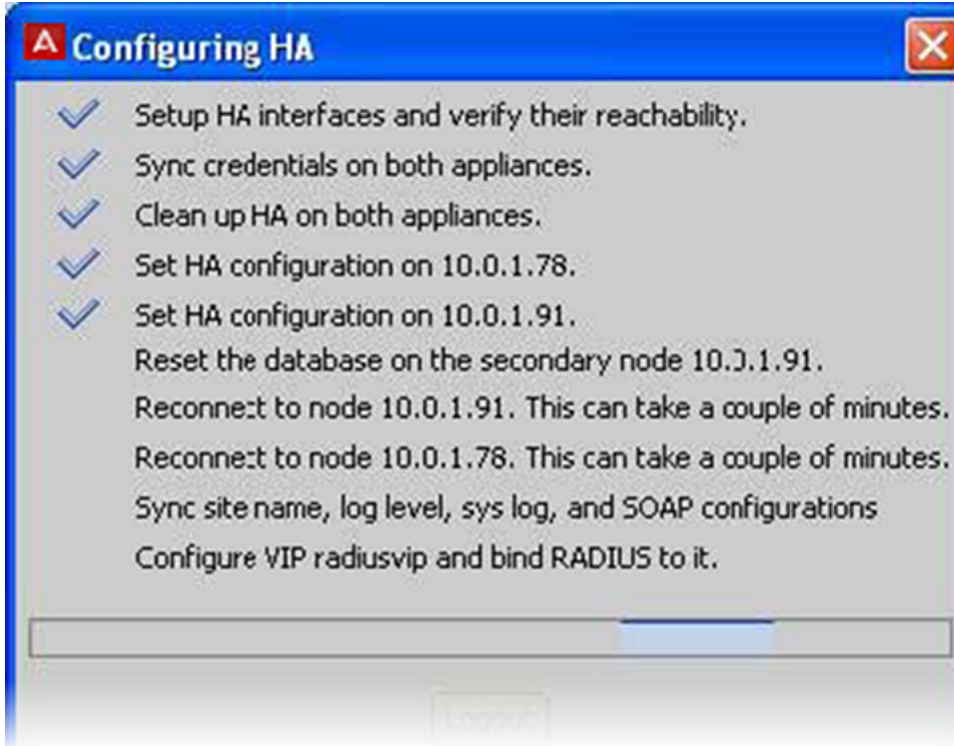
- Click **Next**.
10. The HA Configuration Wizard requires a confirmation that the settings for the two nodes in the new HA pair are correct.



- 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 do not exist, the Wizard displays an error message. For instructions on fixing such problems, see [Problem: HA Set-up fails](#) on page 484 in the Troubleshooting section.

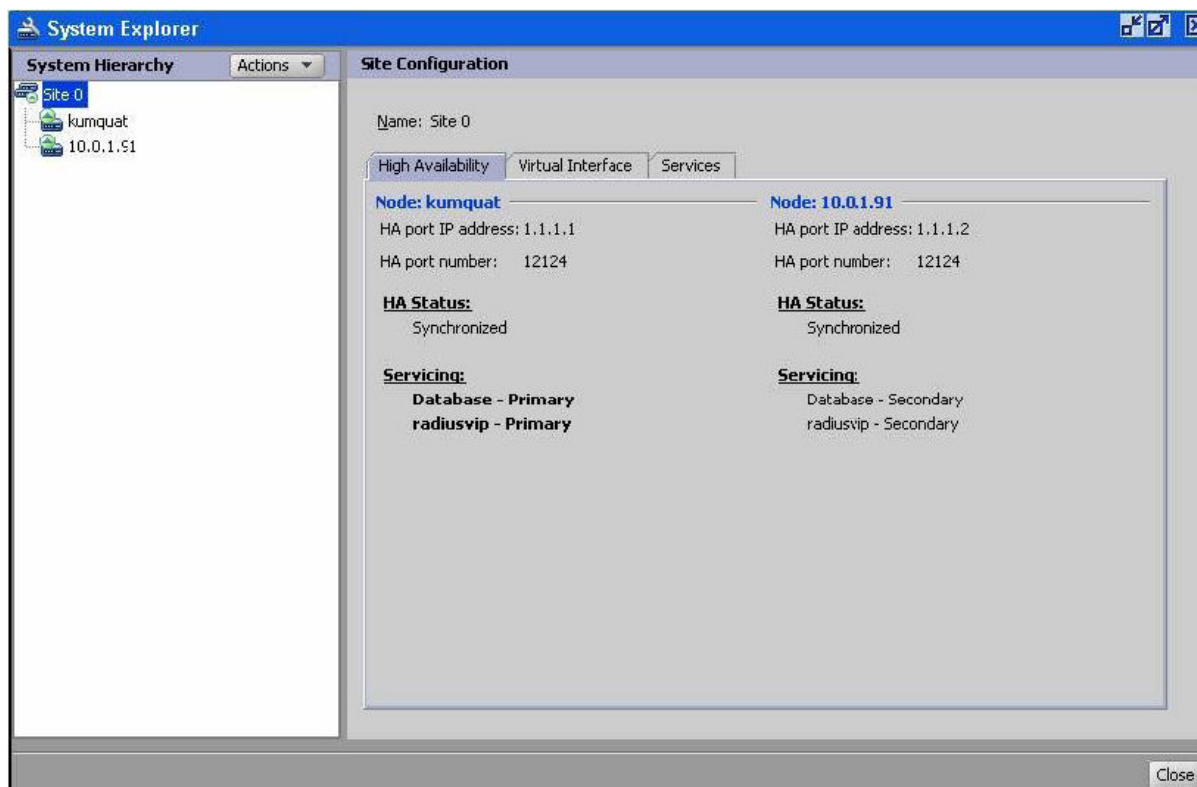
! Important:

If you provide incorrect settings to the HA Configuration Wizard, the wizard's progress bar may appear to freeze. Contact Avaya customer support for assistance. See [Support](#) on page 21.

After setting up the HA link, Dashboard reconnects to the pair.



You manage both Ignition Servers from a single Dashboard session. Ignition Dashboard displays the two nodes successfully linked as a pair.



Bind Services to the VIP

Follow this procedure to bind the RADIUS and/or SOAP services to the VIP interface. This procedure continues from Step 11 in the preceding procedure.

Procedure

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 configured in the Virtual Interface Configuration window in Step 9 in the preceding procedure

6. Click **OK**.

Your Ignition Server HA pair set-up is complete. The next steps are:

- If you have configured the **RADIUS service** on the VIP port, then you must configure 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, see the warning in [Warning for Sites Running Ignition Server in HA Mode](#) on page 204.

- If you have configured the **SOAP API** on the VIP port, then you must configure Ignition Server Guest Manager to connect to the SOAP service at the VIP IP address. See *Configuring Avaya Identity Engines Ignition Guest Manager, NN47280–501* for more information.

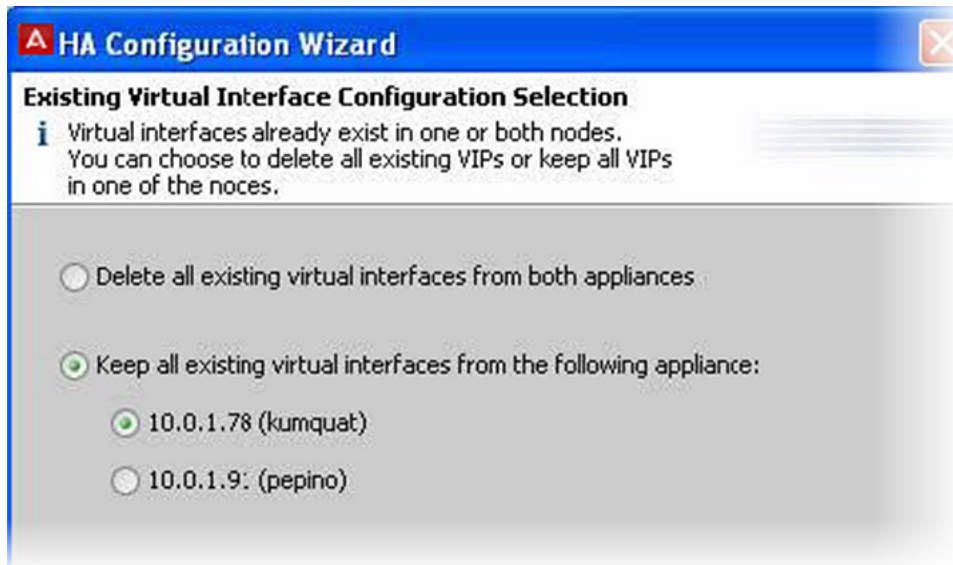
Restoring a saved VIP configuration

When an HA pair is broken through 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.

Use the following procedure to restore a saved VIP configuration.

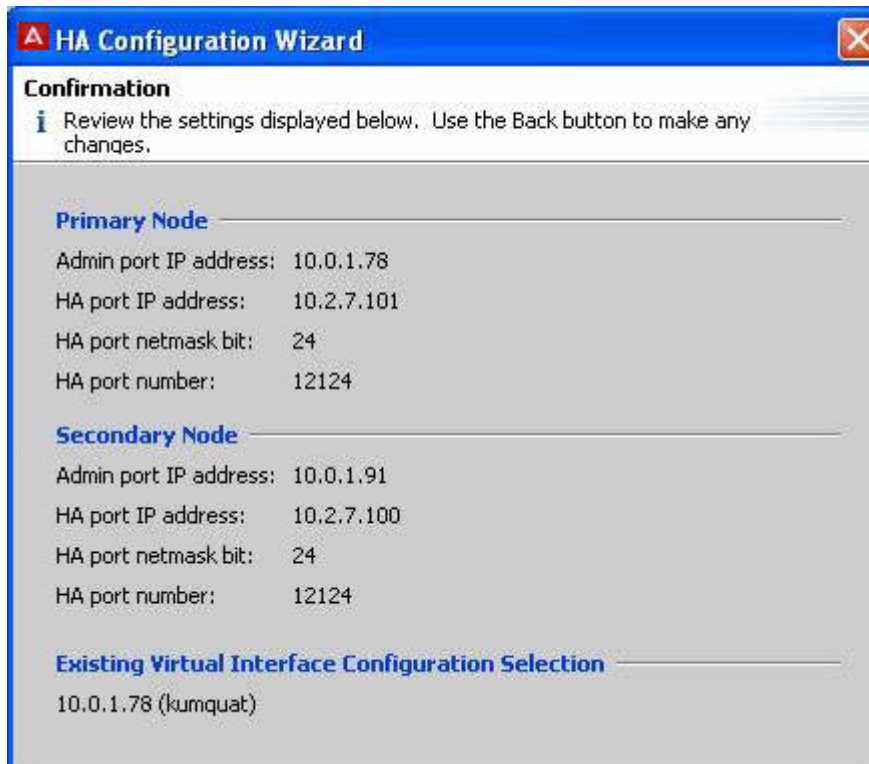
Procedure

1. Run the HA Configuration Wizard, as explained in [Creating an HA Pair](#) on page 392. When Step 8 is completed, the Wizard displays the **Existing Virtual Interface Configuration Selection** window.



2. Select **Keep all existing virtual interfaces...**
3. Select the node whose VIP configuration you want to load.

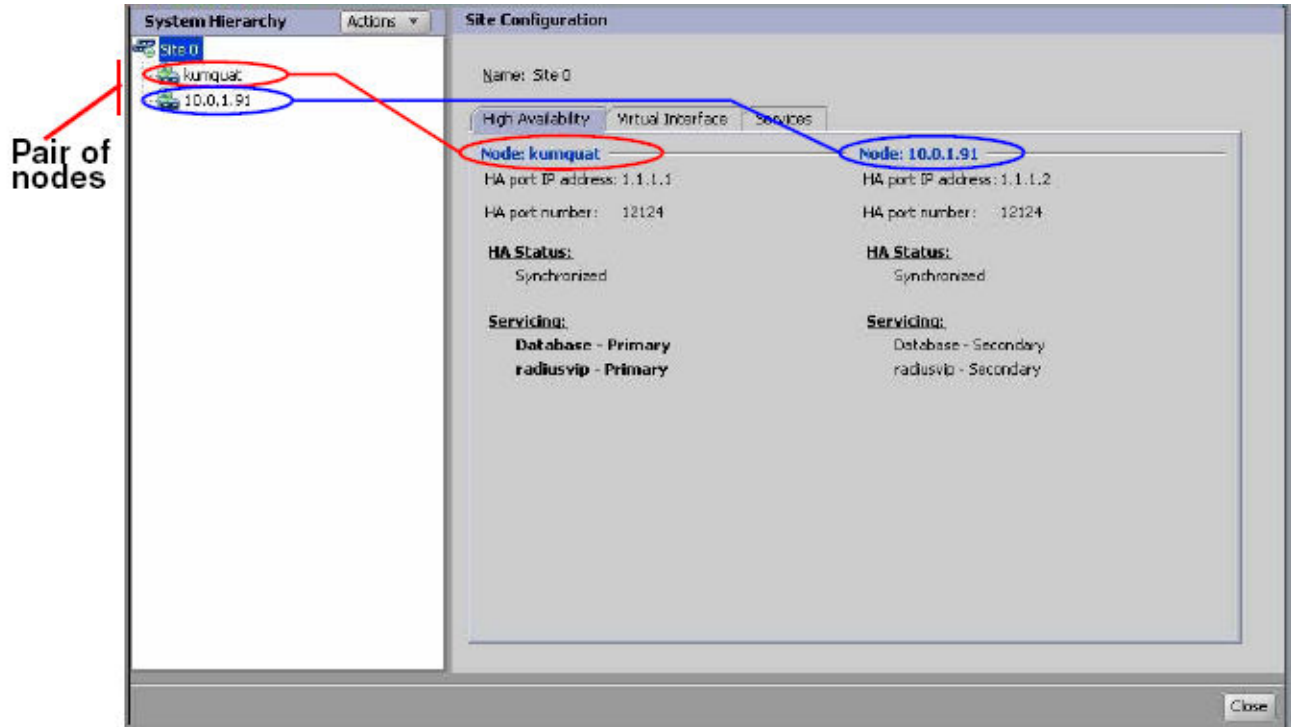
4. Click **Next**.
5. In the **Confirmation** window, review your settings and click **Next** to apply the configuration.



6. Return to Step 1 in [Bind Services to the VIP](#) on page 401 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 following information.

- 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”, “syncing 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 \(VIPs\)](#) on page 405.
- 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 57.

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

Use the following procedure to view the virtual interface settings of your HA pair.

Procedure

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

Use the following procedure to add a virtual interface.

*** Note:**

Do not create more than one VIP on your system.

Procedure

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 **Add** in the **Virtual Interface** tab.
4. Enter the following 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 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:

Procedure

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 405.)
3. Click **Edit**. The selected virtual interface’s settings appear.
4. Edit the fields as needed. For field descriptions, see Step 4 in [AddingVIP](#) on page 406.
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 in the display seen in the Sites panel.

Deleting a VIP

Procedure

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 405) 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 Servers 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](#) on page 409.

Procedure

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**. Alternatively, select **Actions > Break HA Link**.
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.

Procedure

1. Use a console terminal to run the Ignition Server CLI and log in to either of the Ignition Servers 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 second Ignition Server.
4. Run the “ha break” command on the second Ignition Server.

```
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. Use the following procedure to reconnect the HA link.

Procedure

1. Select the site in Dashboard’s Configuration hierarchy tree.
2. Complete the breakage of the pair by selecting **Actions > Break HA Link**.
3. Recreate the HA pair. See [Run the HA Wizard](#) on page 393.

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.

Use the following procedure to reinitialize the nodes that are currently linked as an HA pair.

Procedure

1. Break the link.

See [Breaking an HA pair using Dashboard](#) on page 408.

2. Reinitialize the required node(s) and reconfigure the node(s) as necessary.
See [Reinitializing Ignition Server from Dashboard](#) on page 63.
3. Recreate the HA pair.
See [Run the HA Wizard](#) on page 393.

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.

Use the following procedure to backup Ignition Server data.

Procedure

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 422.

During the backup operation, the pair continues to provide uninterrupted AAA service.

Important:

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.

Next steps

Troubleshooting Backups: If the backup operation on an Ignition Server which belongs to a linked node fails to complete:

- Break the HA link.
- Log in to one of the Ignition Servers in the HA pair.
- Run the backup on this Ignition Server.
- Re-create the required HA pair. See [Run the HA Wizard](#) on page 393.

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 Servers, the **Restore** window displays.

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.

Follow this procedure to restore Ignition Server data.

Procedure

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 425.

Next steps

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.

Important:

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 414.

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 429.

Note:

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 414.

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 Servers 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).

Example:

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.

Assume you want to replace the *venus* Ignition Server.

	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 the Ignition Server to be replaced from production; 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 configure 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 configure 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.
 - Configure the HA interface. (In Dashboard’s Configuration hierarchy tree:
 - Click the Node.
 - Click **Ports: HA Port**.
 - Click **Edit**.

- Enable the port.
 - Configure its IP address.
6. Configure the replacement Ignition Server VM correctly for 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, configure the following settings.
 - Designate the existing production Ignition Server (*saturn* in this example) as Primary.
 - When the Wizard offers the option of deleting or restoring the VIP configuration, click **Keep all existing virtual interfaces**. The Wizard migrates the VIP from the existing/production Ignition Server. (See [Restoring a saved VIP configuration](#) on page 402.) 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 use the following procedure. Instead, follow the steps in [Enabling the service port](#) on page 73.

Example:

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.

Assume you want to change the IP address of the Admin Port interface on *mercury* to 10.0.3.99.

	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
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. 3. Using the CLI, log in to the Ignition Server whose IP address you want to change. Configure its Admin Port IP addresses with the new IP address.

In this example, on *mercury*, use the CLI as follows to configure the IP address.

```
Identity Server> interface admin ipaddr 10.0.3.39/21
```

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 Ignition Server with the new IP address 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 recreate the HA pair. You can initiate the Wizard from either Ignition Server. When the Wizard prompts you for HA configuration information, make the following settings.
 - Designate the still-running production Ignition Server (*neptune* in this example) as Primary.
 - Choose to migrate the VIP from the still-running production Ignition Server. In this example, the Wizard applies the *neptune* VIP settings to the pair and activates the VIP.
 - 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, use the following recovery procedure to recover the boxes.

Procedure

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 `ha break` 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 **Actions > Restore Data**.
3. Connect to Node 2 using **Ignition Dashboard**. Perform a system restore in the same manner as explained in Step 2.
4. While still connected to Node 2, re-create the HA pair: In Dashboard’s hierarchy tree, click the **Site** and select **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 **Keep all existing virtual interfaces**. The Wizard migrates the existing VIP settings to the restored HA Pair. See [Restoring a saved VIP configuration](#) on page 402.


Appendix C: Extended High Availability configuration

Overview

Release 9.1 introduces the first phase of an Extended High Availability (HA) configuration with geographically redundant Avaya Identity Engines (IDE) Ignition servers. One site is designated as a root site, which contains the primary, active Ignition Server. One or more sites are designated as branch sites, which contain secondary, inactive Ignition Servers. Configuration of device and user guest accounts occurs on the root site, and periodic exports of that data transfer the information to the branch sites. In the event of site failure, a branch site can take over access requests.

When you configure an export schedule on an Ignition Server, you designate that server as a root Ignition Server—no further configuration is required. Similarly, when you configure an import schedule on an Ignition Server, you designate that server as a branch server with no further configuration required.

The following diagram shows a typical Extended High Availability configuration. In this example, Data Center 1 (DC1) is the root site, and Data Centers 2, 3, and 4 (DC2, DC3, and DC4) are the branch sites.



Limitations

For Release 9.1, the following limitations apply to Extended high availability configurations:


- Only user and device details are exported to branch Ignition Servers.
- Groups, Access Policies, Internal Provisioners, Authenticators, Directory Services and Directory Sets, and so on, are not exported and must be configured on each individual branch Ignition Server or done through a complete configuration restore into the Ignition Server in the remote Data Center. Then schedule a subsequent export and import of user and device details from the root server to the branch server.
- Internal Provisioners that are created using the Guest Manager administrator application are not exported and must be included in the configuration backup and restore process.
- If user or device details are deleted on the root server, they must be manually deleted on the branch servers as well; import operations do not automatically delete obsolete details on branch servers.
- Device details that are created using the Mobile Device Management (MDM) Directory Service are not included in export or import operations.
- File transfers are executed using the SFTP protocol only.

Configuring scheduled exports

When you configure an export schedule on an Ignition Server, that server is designated as a root server.

Procedure

1. In the Dashboard **Configuration** tree, click the name of your site.
2. Click the **Extended HA** tab.



3. Click **New Export**.

The New Schedule window displays.

The 'New Schedule' dialog box contains the following fields and options:

- Schedule Name:** [Empty text field]
- Schedule Type:** Export
- Optionally Password Protect Export Files:**
 - Password for Export file:** [Empty text field]
 - Confirm Password:** [Empty text field]
- Start Time:** 2015-02-04 07:54:00 [Clock icon]
- Recurrence:**
 - One Time
 - Daily
 - Weekly
 - Monthly
 - Hourly
- Recur Every Week On:**
 - Sunday
 - Monday
 - Tuesday
 - Wednesday
 - Thursday
 - Friday
 - Saturday
- Recurrence Range:**
 - No End Date
 - End By: 2015-02-04 07:54:13 [Clock icon]
- Export Host Settings:**
 - Export To Host:** [Empty text field]
 - LogIn Name:** [Empty text field]
 - Password:** [Empty text field]
 - Confirm Password:** [Empty text field]
 - Number of Export Files:** 1
 - Destination Path:** [Empty text field]
 - Export Filename:** [Empty text field]

Buttons: OK, Cancel

4. Type a **Schedule Name** to identify the schedule.
5. In the **Password for Export file** and **Confirm Password** fields, enter a password for encrypting the export file. Leave these fields blank if you do not want to encrypt the file. When you import this file on the branch node, you must supply the same password to decrypt the file.
6. Select the start time and schedule.
 - Click the clock and calendar icon of the **Start Time** field to set the time when the first export will begin.
 - In the **Recurrence** section, specify the frequency (one time, daily, weekly, monthly, or hourly).
 - If applicable, to the right of the **Recurrence** field, specify the detailed frequency parameters as applicable. For monthly, choose a numbered day of the month. For weekly,

choose a day of the week. For daily, specify either a frequency (“**Every n days**”) or choose **Every Weekday**. To export a file every day, specify a frequency of “Every 1 days”.

- In the **Recurrence Range** field, specify the end date, if any, for this schedule.

7. In the **Export Host Settings** section, specify the SFTP server that is to receive export files:
 - In the **Export to Host** field, specify the machine name or IP address of the destination SFTP server.
 - In **Login Name** and **Password**, enter the user name and password of the SFTP server account that will own the export files.
 - Type the password again in the **Confirm Password** field to confirm it.
8. Specify where the file should be exported to, and how many files should be kept on the file server:
 - In the **Destination Path** field, specify the path where the exports are saved on the SFTP server.
 - In the **Number of Export files** field, specify the maximum number of export files to be kept. For example, if you run weekly exports and set this to “3”, then Ignition Server saves three export files over the first three weeks, and in each subsequent week it overwrites the oldest export file.
 - In the **Export Filename** field, specify a name for the export file.
9. 8. Click **OK**.
Your schedule has been saved.
10. To add more export schedules, repeat this procedure.


Configuring scheduled imports

When you configure an import schedule on an Ignition Server, that server is designated as a branch server.

The recommended configuration for Extended high availability uses one file server for export and import operations. If you move the exported file to a secondary file server, ensure that you specify the correct secondary file server information in this procedure.

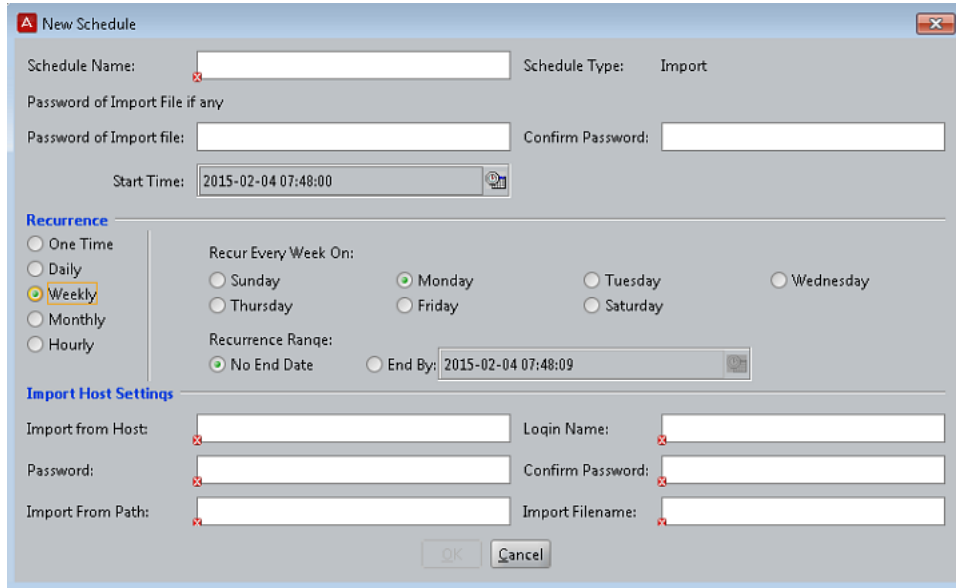
Procedure

1. In the Dashboard **Configuration** tree, click the name of your site.
2. Click the **Extended HA** tab.



3. Click **New Import**.

The New Schedule window displays.



4. Type a **Schedule Name** to identify the schedule.

5. In the **Password of Import file** and **Confirm Password** fields, enter a password for decrypting the import file. Leave these fields blank if the export file was not encrypted. If the export file was encrypted, you must enter the same password that was used to encrypt the file when configuring the export schedule.

6. Select the start time and schedule:

- Click the clock and calendar icon of the **Start Time** field to set the time when the first import will begin.
- In the **Recurrence** section, specify the frequency (one time, daily, weekly, monthly, or hourly).
- If applicable, to the right of the **Recurrence** field, specify the detailed frequency parameters as applicable. For monthly, choose a numbered day of the month; for weekly, choose a day of the week; for daily, specify either a frequency (“**Every n days**”) or choose **Every Weekday**. To import a file every day, specify a frequency of “Every 1 days”.
- In the **Recurrence Range** field, specify the end date, if any, for this schedule.

7. In the **Import Host Settings** section, specify the SFTP server that stores the file specified in the export schedule.:


- In the **Import from Host** field, specify the machine name or IP address of the SFTP server. If you are using the recommended configuration of one file server for both export and import operations, this must match the information specified for the export schedule.
- In **Login Name** and **Password**, enter the user name and password of the SFTP server account that stores the export files.
- Type the password again in the **Confirm Password** field to confirm it.

8. Specify where the file should be imported from:
 - In the **Destination Path** field, specify the path where the export files are saved on the SFTP server. If you are using the recommended configuration of one file server for both export and import operations, this must match the destination specified for the export schedule.
 - In the **Import Filename** field, enter the name of the file, which must match the filename specified in the export schedule.
9. 8. Click **OK**.
Your schedule has been saved.
10. To add more import schedules, repeat this procedure.

Editing an export or import schedule

Procedure

1. In the Dashboard **Configuration** tree, click the name of your site.
2. Click the **Extended HA** tab.
The Extended HA tab displays, listing all existing import and export schedules.
3. Highlight the desired schedule.



4. To delete a schedule, click **Delete** and click **OK** to confirm.
5. To edit a schedule, click **Edit**, make the desired changes, and click **OK**.

Appendix D: Backup and Restore Procedures

This appendix explains how to back up and restore the Avaya Identity Engines Ignition Server data and configuration.

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:

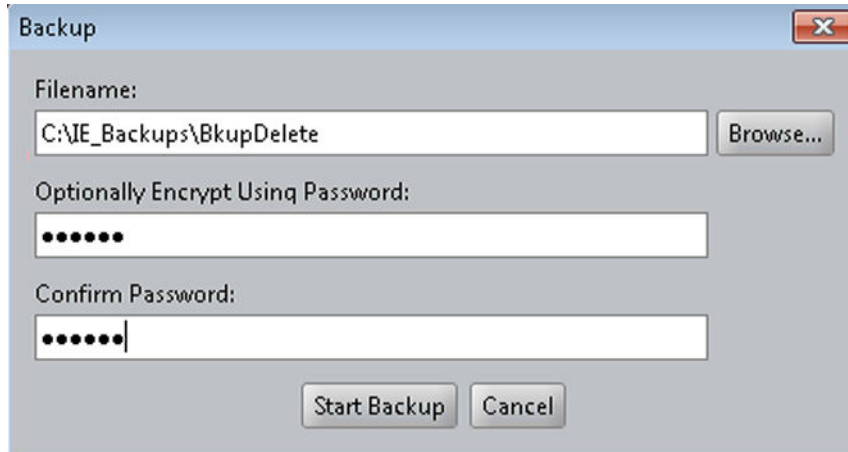
- policies (tunneling, identity routing, authentication, and authorization)
- users and groups in the Ignition Server internal store
- authenticator/NAS configuration
- certificates

Creating a backup

Follow this procedure to create a backup of the data on your Ignition Server.

Procedure

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**.
 - Click on the site and then choose **Actions > Backup Data**.
3. Dashboard displays the **Backup** window.



4. Click **Browse** to specify the path and filename for the backup file.

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 9.x system with a backup file taken from a 8.x system or 9.x system), it is safe practice to note the Ignition Server firmware version in the file name of your backup.

5. In the **Optionally Encrypt Using Password** and **Confirm Password** fields, enter a password 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 password 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.

Configuring scheduled backups

Your Ignition Server or HA pair of Ignition Servers 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.

Procedure

1. In the Dashboard **Configuration** tree, click the name of your site.
2. Click the **Scheduled Backups** tab.



3. Click **New Backup**.

The New Schedule window displays.

4. Type a **Schedule Name** to identify the schedule.
5. In the **Password for Backup** and **Confirm Password** fields, enter a password for encrypting the backup 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 password to decrypt the file.
6. Select 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).
 - If applicable, to the right of the **Recurrence** field, specify the detailed frequency parameters. For monthly, choose a numbered day of the month; for weekly, choose a day

of the week; for daily, 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.
7. In the **Backup Host Settings** section, 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.
 - In **Login Name** and **Password**, enter 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.
 8. Specify where the backup file should be exported to, and how many files should be kept on the file 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 saves three backup files over the first three weeks, and in each subsequent week it overwrites the oldest backup file.
 - In the **Destination Path** field, specify the path where the backups are saved on the SFTP server.
 - In the **Backup Filename** field, specify a name for the backup file.
 9. 8. Click **OK**.
- Your schedule has been saved.
10. If you wish to add more backup schedules, repeat this procedure.

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 9.x system with a backup file taken from 8.x system or 9.x 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 *Avaya Identity Engines 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 to 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.


Restoring data from a backup file

Follow this procedure to restore your Ignition Server from a backup file.

Procedure

1. Open the Configuration view of Dashboard.
2. In the **Configuration Hierarchy** panel, right-click on the **Site**. Select **Restore Data**.

The **Restore** window displays.



3. Click **Browse** to specify the path and filename for the backup file from which you are restoring the data.
4. In the **Password** field, enter the password that you entered when encrypting the backup file. Leave the field blank if the file is not encrypted.
5. Select the types of system configuration data that you want to restore. Select as many as required.
 - **Primary node network configuration:** The network settings specific to the primary node (see [Managing a node](#) on page 62).

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 check box restores the records of Ignition's internal data store, including policies and directory service settings.
- **License configuration:** *All of the Ignition Server licenses* that were on the backed up machine.

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 can click **Logout** 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 primary node network 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 are restoring. In the case of change of Admin IP due to restore (backed up config), Dashboard does not reconnect. In this case, reconnect by entering the new Admin IP address in the Dashboard login window. This information is displayed in the restore window.
9. **Ignition Server licenses:** If you chose to restore the license 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 are 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 78.

Troubleshooting

If, for any reason, the restore operation fails, Ignition Server issues one of the following error messages.

Error Message	Steps to Correct
This is not a valid backup	<p>You have assigned the wrong backup file. Use the browser on your personal computer or workstation and locate the correct backup file.</p> <p>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</p>

Table continues...

Backup and Restore Procedures

Error Message	Steps to Correct
	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 password you provided for encryption when you created the backup file.

Appendix E: 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.

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 429
- [View the current Firmware version and all installed Images and Packages](#) on page 430

To determine the version of Ignition Dashboard you are running, select **Help > About** from the main window.

View the current Firmware version

View the current firmware version on the Ignition Server.

Procedure

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.

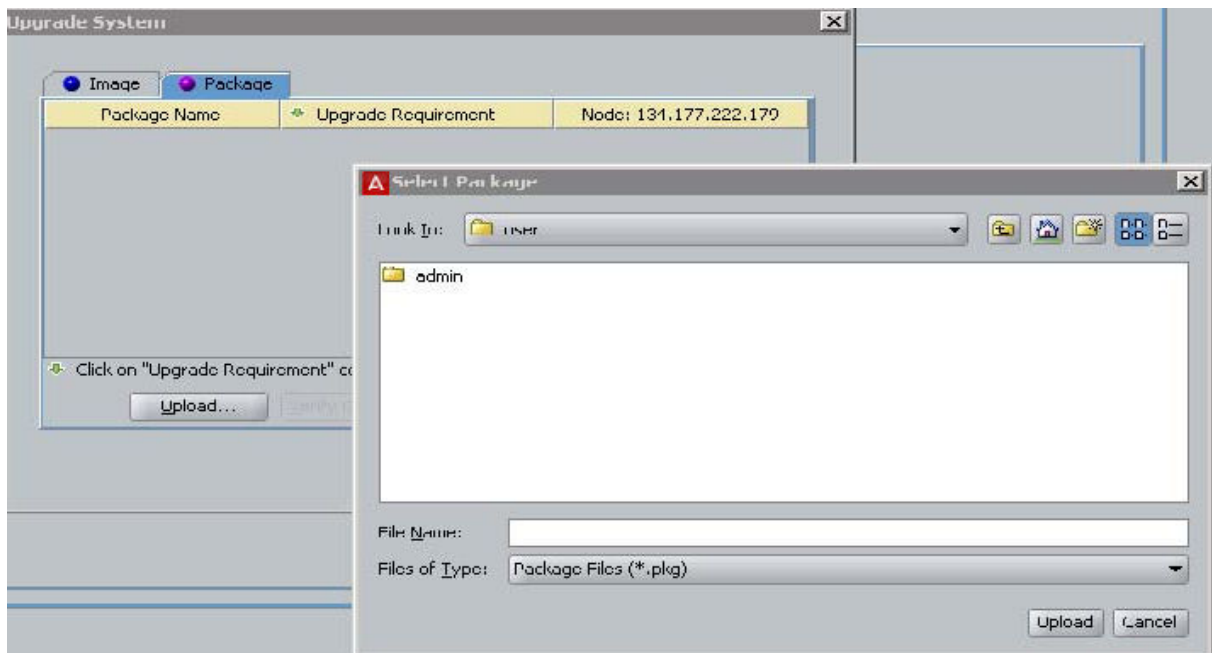
View the current Firmware version and all installed Images and Packages

Follow this procedure to view the current firmware version and all installed firmware images on the Ignition Server.

Procedure

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**.



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 the same. 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 the image. To migrate your Ignition Server to one of these files, see [Activating a firmware image or package](#) on page 432.

The version number takes the form, “LINUX-VM_09_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, “09_00_00” indicates this is firmware version 9.0.0, 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.

Procedure

1. Select a firmware update from the Avaya web site (see [Support](#) 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 **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**.

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**.

* 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 9.0 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 434.

7. Configure the Ignition Server to use the new firmware. See [Activating a firmware image or package](#) on page 432.

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, use the following procedure to activate the firmware image.

Procedure

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**.
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.



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.


Activating the firmware package

Follow this procedure to activate the firmware package.

Only *not-installed* packages can be deleted or verified or activated. Each package contains a flag that indicates when it was installed. For not-installed packages, the Node column displays **Not activated**. Once the package is activated, the display changes to “Activated on MM/DD/YYYY”.

Procedure

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 **Actions > Upgrade System**.
3. Select the **Packages** tab.
4. In the **Firmware Manager** window, select the package you want to activate.
5. In the **Upgrade Requirement** column, select one of the available upgrade requirement/recommendations for the system/server restart. The recommended option is tagged with the check mark. You can select the other options, but Avaya recommends that you should always select the recommended option.




The Upgrade Requirement column contains the following values.

- OS Reboot - 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 about the 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 Servers. The **Activate** button is enabled only when you select an image that is installed on both Ignition Servers, as shown in the following figure.



When you click **Activate**, the **Rebooting** window that indicates the progress of the operation is display 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.

Delete a firmware file.

Procedure

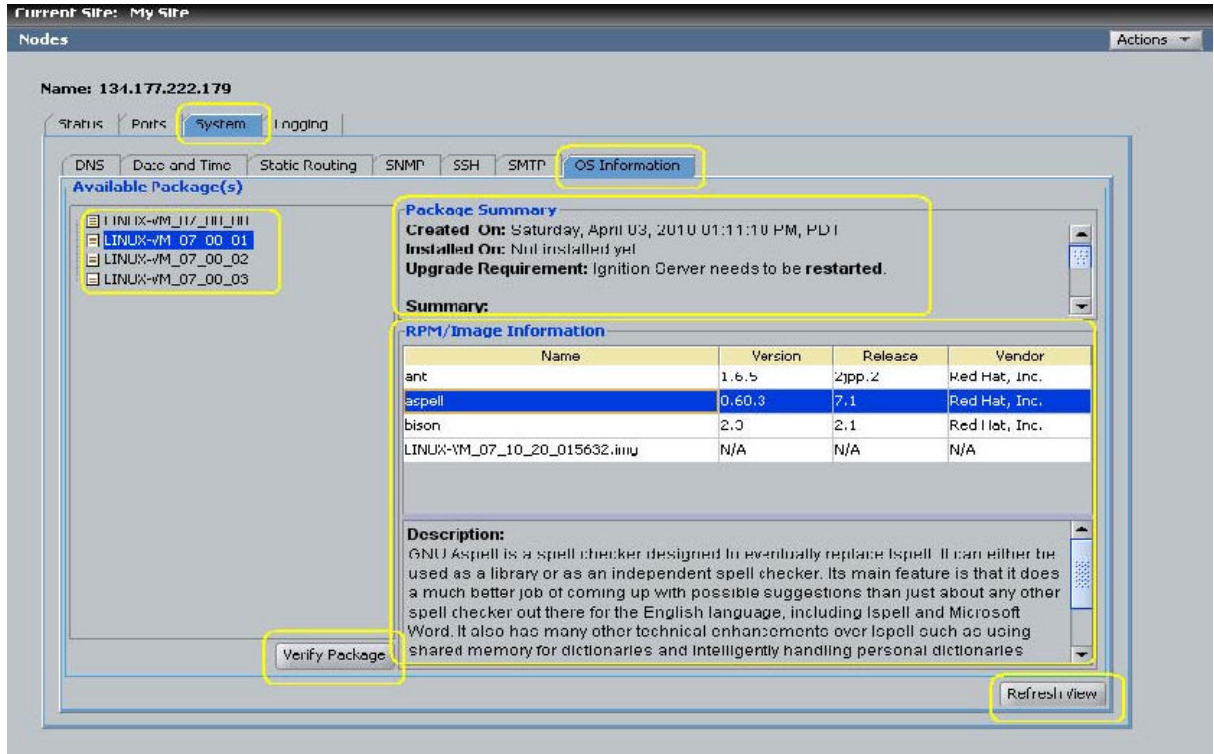
1. In Dashboard's Configuration hierarchy, select your Site.
2. Click **Actions**. (Alternatively, right-click the **Site** in the **Configuration Hierarchy** tree.)
Select **Upgrade System**.
3. Select the **Images** or **Packages** tab.
4. The **Firmware Manager** window appears displaying the existing firmware images or packages on your Ignition Server.
The running version is marked with an asterisk.
5. Scan the list for the firmware file you no longer want to keep, and click on the row to select it.
6. Click **Delete**.
A confirmation window appears.
7. Click **OK** to confirm.
Ignition Server deletes the firmware image.

Viewing Image and Package information

Follow this procedure to view OS information currently installed or available for upgrade

Procedure

1. Launch the **OS information** window by selecting **Nodes > System**.
2. Select the **OS information** tab.



The OS Information window displays the following Read-Only information.

- **Package Summary:** displays installation information.
- **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.
- **Description:** image or file details as a description.

Click **Verify Package** to verify/validate the package content for the package(s) that have not been installed. This button is disabled when the package which is already installed.

You can refresh the data by clicking **Refresh View**.

Upgrading Ignition Server

Ignition Server does not support package Rollback.

! Important:

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 Ignition Server

For an Ignition Server upgrade, there is no roll back option. Avaya recommends using the following procedure to create a backup in the event you need to revert back to a previous Ignition Server release.

Procedure

1. Export and save your existing Ignition Server configuration.
2. Halt the Ignition Server Nodes you plan on upgrading.
3. Make sure the nodes are in a powered-down state.
4. Using the Storage management section of your ESXi Server, create a separate folder to hold the disk drive contents 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 beginning the upgrade.

6. Upload the new release image and activate the image.
The system reboots twice in order to perform the upgrade. At the end of the second reboot, the new release is running.
7. If the upgrade fails or you decide to revert back to the previous release, perform the following steps.
 - a. Power off the VM as the disk is not in a usable state.
 - b. Go to the VM Setup menu and delete the hard disk from the VM. Do this for each node in your system.
 - c. Copy the saved VM back to its original location.
 - d. Attach the disk to VM using the **attach existing disk** option. Do this for each node in your system.
 - e. Power on the VM(s).
 - f. Your previous Ignition Server environment is back online.

Appendix F: 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.

Setting User Preferences

You can set your log viewing preferences in the **Preferences** Window. See [Configuring administration preferences](#) on page 51.

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.

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 438
 - [Setting the Level of Logging to be recorded](#) on page 440
 - [Setting up FTP log export](#) on page 440
-

Turning on logging

Procedure

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 Levels window, activate logging for the log types that 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 *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 **OK**.

Setting the Level of Logging to be recorded

See the section, [Turning on logging](#) on page 438.

Setting up FTP log export

Follow the steps below to set up FTP exporting of your log files.

Procedure

1. In Ignition Dashboard, click Configuration 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**.

Type	Auto Export When Capacity Reached	Export Periodically	Start Periodic Export
Access	<input checked="" type="checkbox"/>	Daily ▾	2015-03-10 10:17:41
Audit	<input checked="" type="checkbox"/>	Daily ▾	2015-03-10 10:17:44
Security	<input checked="" type="checkbox"/>	Daily ▾	2015-03-10 10:17:48
System	<input checked="" type="checkbox"/>	Hourly ▾	2015-03-10 10:17:51
Access Details	<input checked="" type="checkbox"/>	Weekly ▾	2015-03-10 10:17:55

Log Export Host Settings

Export To Host: Login Name:

Password: Confirm Password:

OK Cancel

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 click 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.
 - Click **OK**.

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.

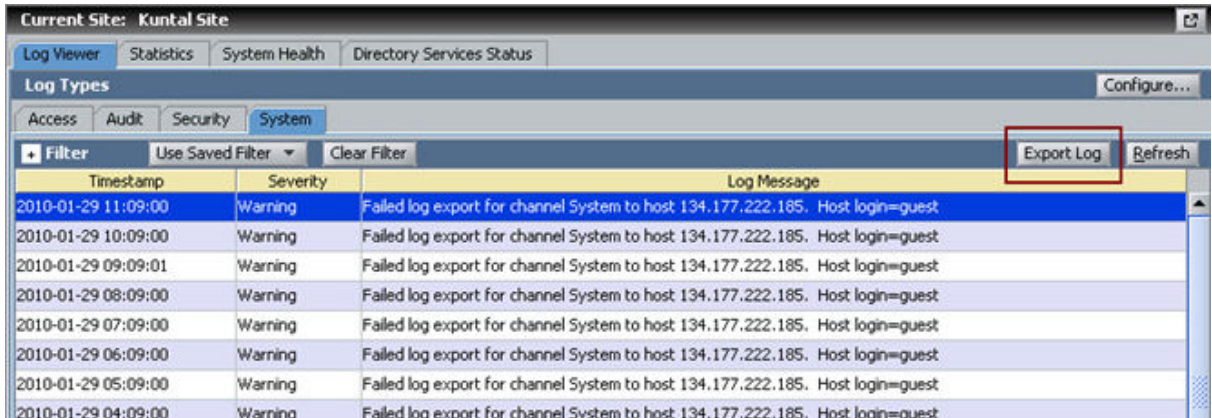
Exporting logs using the Export All variable may take longer depending on the number of logs that are in the database.

Follow the steps below to export your log files immediately:

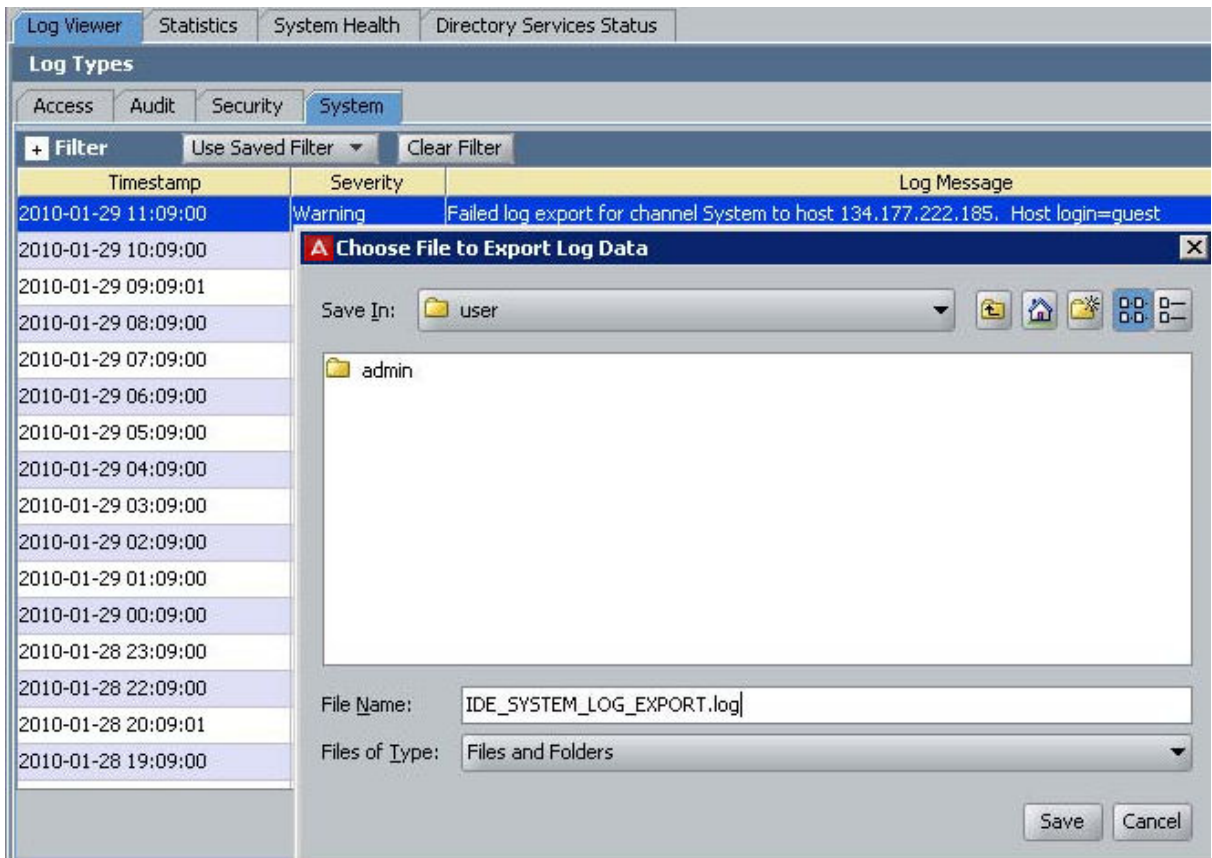
Procedure

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 then click on the **Actions** drop down box to enable the **Debug Logs** and **Advanced Log Levels** options.
5. Highlight the logs you want exported.

Setting up logging



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.



8. Click **Save**.
9. You can use any text editor such as Textpad, Notepad, or vi, to open the exported file.

```

-----
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 ***
-----

```

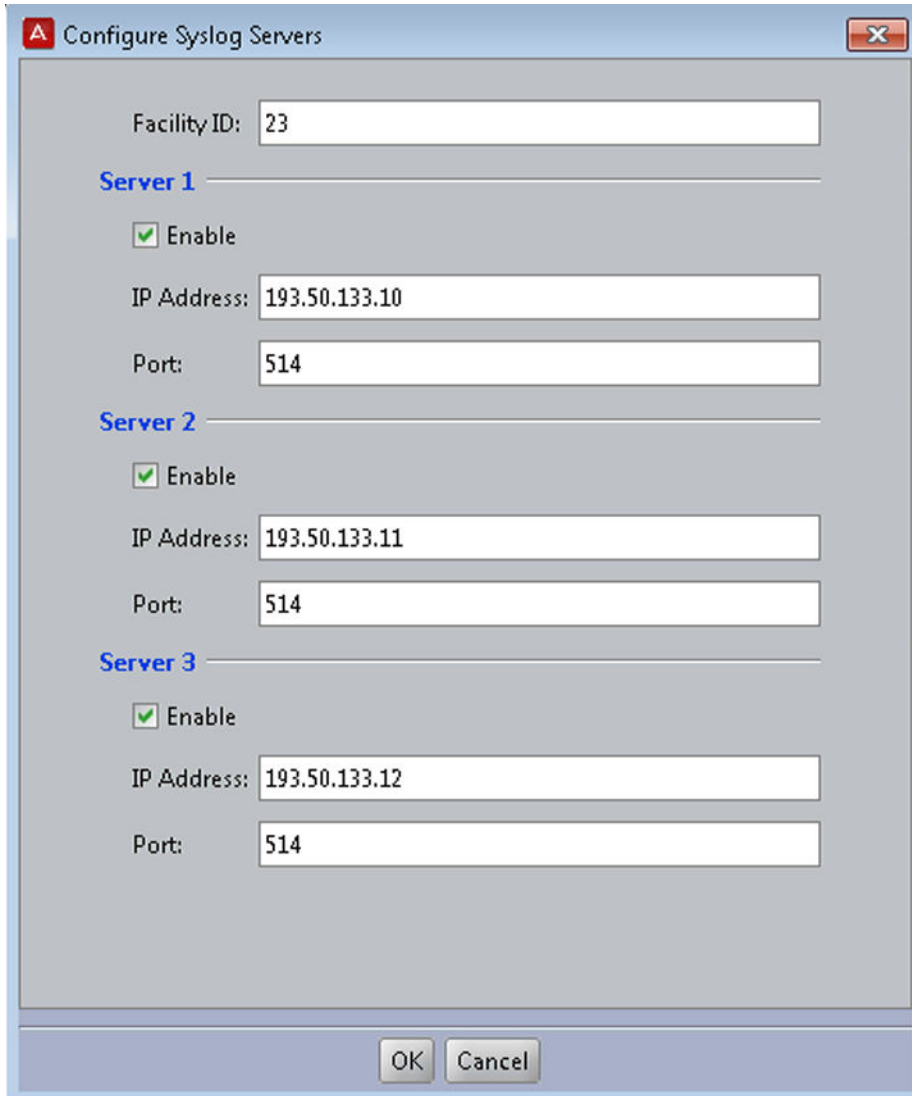
Timestamp	Severity	Log Message
2010-01-29 11:09:00	Warning	Failed log export for channel System to h
2010-01-29 10:09:00	Warning	Failed log export for channel System to h
2010-01-29 09:09:01	Warning	Failed log export for channel System to h
2010-01-29 08:09:00	Warning	Failed log export for channel System 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

Direct log messages to one or more Syslog servers.

Procedure

1. In Ignition Dashboard, click **Configuration** 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**.
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 **OK**.

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.

Procedure

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. 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.
 - 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 through 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.

Procedure

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.

Edit SNMP Configuration

Please provide SNMP Configuration

Enabled:

UDP Port: 161

Bound Interface: Admin Port

Community String: hBk580VA

Location: Sunnyvale-Lab-09

Contact Address: network-admin@zipwaves.com

Source IP Addresses:

IP address	Netmask
204.198.100.0	24

Add

Delete

OK Cancel

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.

Procedure

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_009734S
```

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 HOSTRESOURCES-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, ipRoutelfIndex, 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.

Interface Name	Interface Name in Ignition Server firmware/CLI	Index number in SNMP records	SNMP Example
Loopback Address	lo	1	ifDescr.1
Admin port	eth0	2	ifDescr.2
Service Port	eth1	3	ifDescr.3
HA port	eth2	4	ifDescr.4

Appendix G: Viewing logs and statistics

This appendix 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.

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+, Guest Manager, SAML, and Administration). See “[Access Log: RADIUS and TACACS+ Accounting](#)” on page 454 and [AAA Summary tabs](#) on page 469.
- Highly detailed information about a single login attempt and its results. See [Access Record Details](#) on page 456 .
- List of current logged-in users. See [User Accounting tab](#) on page 471.
- List of current AD-authenticated devices. See [Learned Devices tab](#) on page 473.
- Audit log of administrator actions on the Ignition Server. See [Audit Log](#) on page 460.
- Security and system health logs. See [System Health tab](#) on page 468.
- Statistics and logs detailing authentication/authorization transactions, including statistics categorized by authentication protocol. See [Statistics tab](#) on page 464.
- Directory service interaction statistics and logs. See [Directory Services Status Tab](#) on page 468.

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 438.)

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

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

Procedure

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.

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](#) on page 451.)

Filtering your view of the Logs

Follow this procedure to apply a filter.

Procedure

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 and 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 452 explains the **Filtering Criteria**.
5. Click **Apply**.

Only records matching all your criteria are shown.


Timestamp	Severity	Log Message
2010-01-29 11:09:00	Warning	Failed log export for channel System to host 134.177.222.185. Host login=gu
2010-01-29 10:09:00	Warning	Failed log export for channel System to host 134.177.222.185. Host login=gu
2010-01-29 09:09:01	Warning	Failed log export for channel System to host 134.177.222.185. Host login=gu
2010-01-29 08:09:00	Warning	Failed log export for channel System to host 134.177.222.185. Host login=gu
2010-01-29 07:09:00	Warning	Failed log export for channel System to host 134.177.222.185. Host login=gu
2010-01-29 06:09:00	Warning	Failed log export for channel System to host 134.177.222.185. Host login=gu
2010-01-29 05:09:00	Warning	Failed log export for channel System to host 134.177.222.185. Host login=gu
2010-01-29 04:09:00	Warning	Failed log export for channel System to host 134.177.222.185. Host login=gu

Criteria for filtering Log messages

The following table lists the filter criteria and available test values for each in the logging tabs. Not all criteria are available in all tabs.

Command or Button	Purpose
Date and Time Period	Choose one of the following. <ul style="list-style-type: none"> • 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 <i>System</i> log only. Choose a Log Level to display only messages of the severity level you select. From least severe to most severe, the

Table continues...

Command or Button	Purpose
	<p>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.</p> <p> Warning:</p> <p>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.</p>

The following table explains the buttons you use to add, remove, apply, and manage filters.

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.

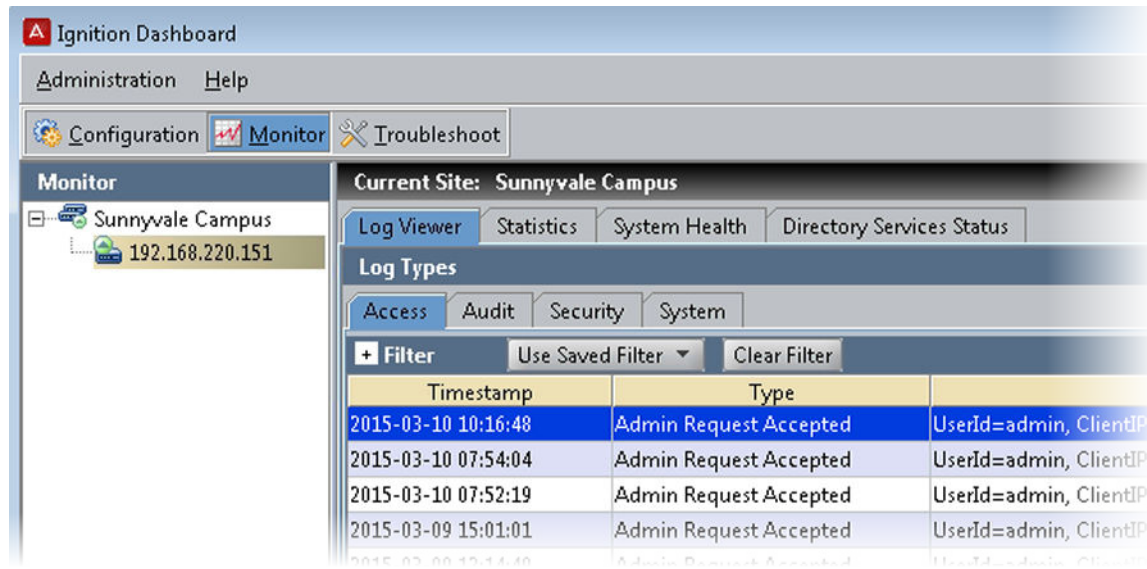
Managing the stored logs

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 440.

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.



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

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


View the Access log.

Procedure

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 456.
5. You can filter the set of records. See [Filtering your view of the Logs](#) on page 451.

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.



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.

Entry Name	Description
acct-status-type	Event type, which is one of (START, STOP, or UPDATE)

Table continues...

Entry Name	Description
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 58. 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 440.

Access Record Details

The Access Record Details window shows the submitted details and returned results of a user's or device's login attempt.

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 52.

The following example shows Access Record Details displayed in a dedicated panel.

Current Site: Sunnyvale Campus

Log Viewer | Statistics | System Health | Directory Services Status

Log Types

Access | Audit | Security | System

+ Filter | Use Saved Filter | Clear Filter

Timestamp	Type	
2015-03-10 10:16:48	Admin Request Accepted	UserId=admin, Clie
2015-03-10 07:54:04	Admin Request Accepted	UserId=admin, Clie
2015-03-10 07:52:19	Admin Request Accepted	UserId=admin, Clie
2015-03-09 15:01:01	Admin Request Accepted	UserId=admin, Clie
2015-03-09 12:14:49	Admin Request Accepted	UserId=admin, Clie
2015-03-09 08:51:41	Admin Request Accepted	UserId=admin, Clie
2015-03-09 08:05:17	Admin Request Accepted	UserId=admin, Clie
2015-03-09 06:51:33	Admin Request Accepted	UserId=admin, Clie
2015-03-09 06:35:59	Admin Request Accepted	UserId=admin, Clie
2015-03-05 12:04:01	Admin Request Accepted	UserId=admin, Clie
2015-03-05 08:49:19	Admin Request Accepted	UserId=admin, Clie
2015-03-03 14:31:40	Admin Request Accepted	UserId=admin, Clie
2015-02-26 22:45:33	Admin Request Accepted	UserId=admin, Clie
2015-02-26 16:42:50	Admin Request Accepted	UserId=admin, Clie

[Access Record Details...](#)

Access Result: Allow

Authentication Details

Access Policy:

ID: admin

Client IP Address: 192.168.220.161

Lookup Service:

Authentication Service:

Click to display the Access Record Detail in its own window.

The following example shows Access Record Details displayed as a tooltip.

Current Site: Sunnyvale Campus

Log Viewer | Statistics | System Health | Directory Services Status

Log Types

Access | Audit | Security | System

+ Filter | Use Saved Filter | Clear Filter

Timestamp	Type	
2015-03-10 10:16:48	Admin Request Accepted	UserId=admin, ClientIP=192.168.220.161,
2015-03-10 07:54:04	Admin Request Accepted	UserId=admin, ClientIP=192.168.220.161,
2015-03-10 07:52:19	Admin Request Accepted	UserId=admin, ClientIP=192.168.220.161,
2015-03-09 15:01:01	Admin Request Accepted	UserId=admin, ClientIP=192.168.220.160,
2015-03-09 12:14:49	Admin Request Accepted	UserId=admin, ClientIP=192.168.220.160,
2015-03-09 08:51:41	Admin Request Accepted	92.168.220.160,
2015-03-09 08:05:17	Admin Request Accepted	92.168.220.161,
2015-03-09 06:51:33	Admin Request Accepted	92.168.220.160,
2015-03-09 06:35:59	Admin Request Accepted	92.168.220.161,
2015-03-05 12:04:01	Admin Request Accepted	92.168.220.160,
2015-03-05 08:49:19	Admin Request Accepted	92.168.220.160,
2015-03-03 14:31:40	Admin Request Accepted	92.168.220.160,
2015-02-26 22:45:33	Admin Request Accepted	92.168.220.160,
2015-02-26 16:42:59	Admin Request Accepted	92.168.220.160,
2015-02-26 13:56:29	Admin Request Accepted	92.168.220.160,
2015-02-25 16:56:54	Admin Request Accepted	92.168.220.160,
2015-02-25 15:02:26	Admin Request Accepted	92.168.220.160,
2015-02-25 12:18:58	Admin Request Accepted	UserId=admin, ClientIP=192.168.220.160,
2015-02-25 10:51:54	Admin Request Accepted	UserId=admin, ClientIP=192.168.220.160,

Access Result: Allow

Authentication Details

Access Policy:

ID: admin

Client IP Address: 192.168.220.160

Lookup Service:

Authentication Service:

Decision: Authenticated

Authorization Details

Policy Rule Used:

Decision: Allow

Role: sys-admin

Viewing the Access Record details

Procedure


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 displays.

! Important:

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 on your computer's clipboard as text. You can paste them into any text editor or word processor.



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, enablepassword-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.

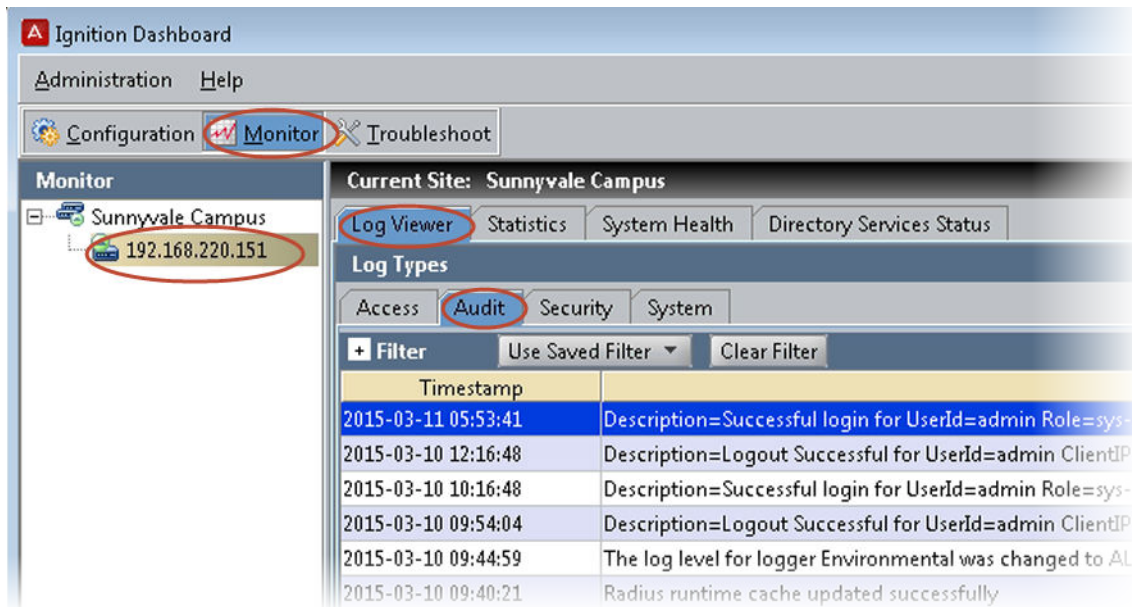
Activating the recording of Access Record Details

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 **Configure** in the upper-right corner of the window.
5. In the **Configure Log Types** window, select the **Access Details: Enabled** check box to turn on detailed logging, or clear the check box to turn it off.

Audit Log

The audit log records administrative actions done on the Ignition Server.



Contents of 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

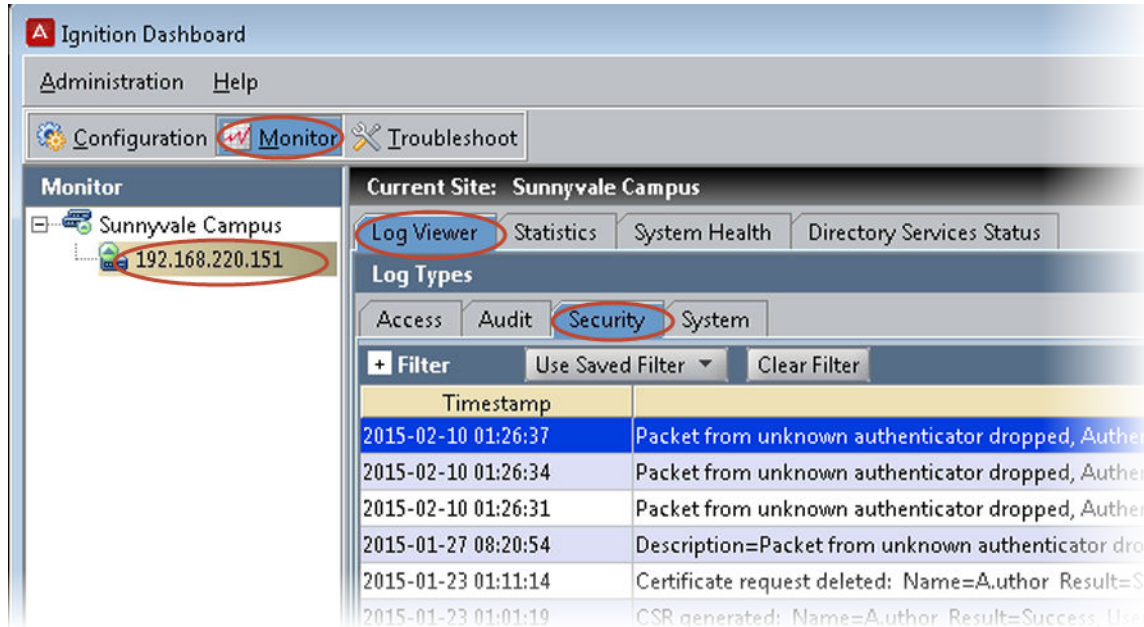
Viewing the Audit Log

Procedure

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 451.

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

Procedure

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 451.

System Log

The following example shows the System Log.



Contents of the System Log

The System log contains 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.

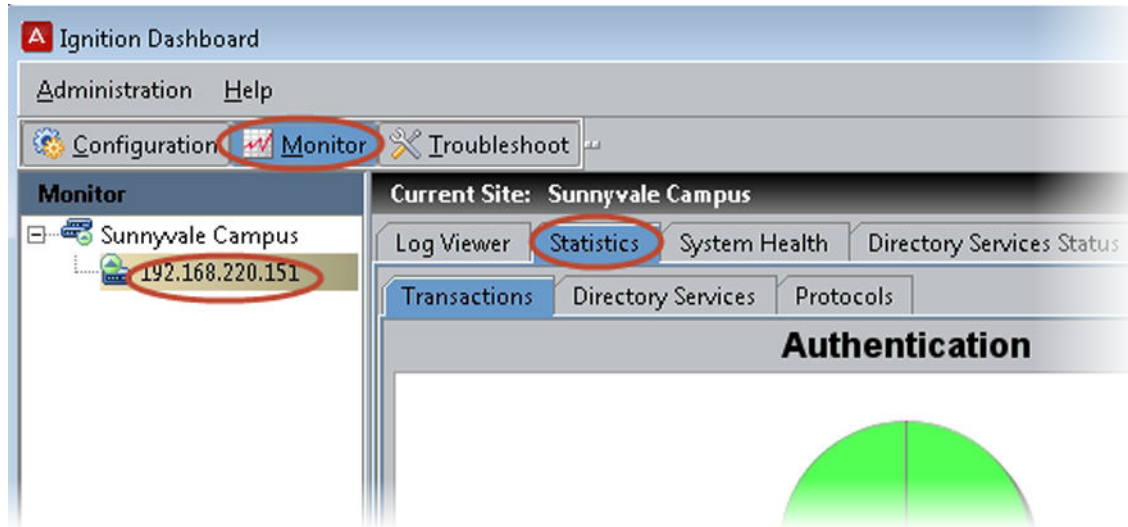
To minimize the number of System logging events recorded, log only events with a severity level of **Fatal**.

Viewing the System Log

Procedure

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 451.

Statistics tab

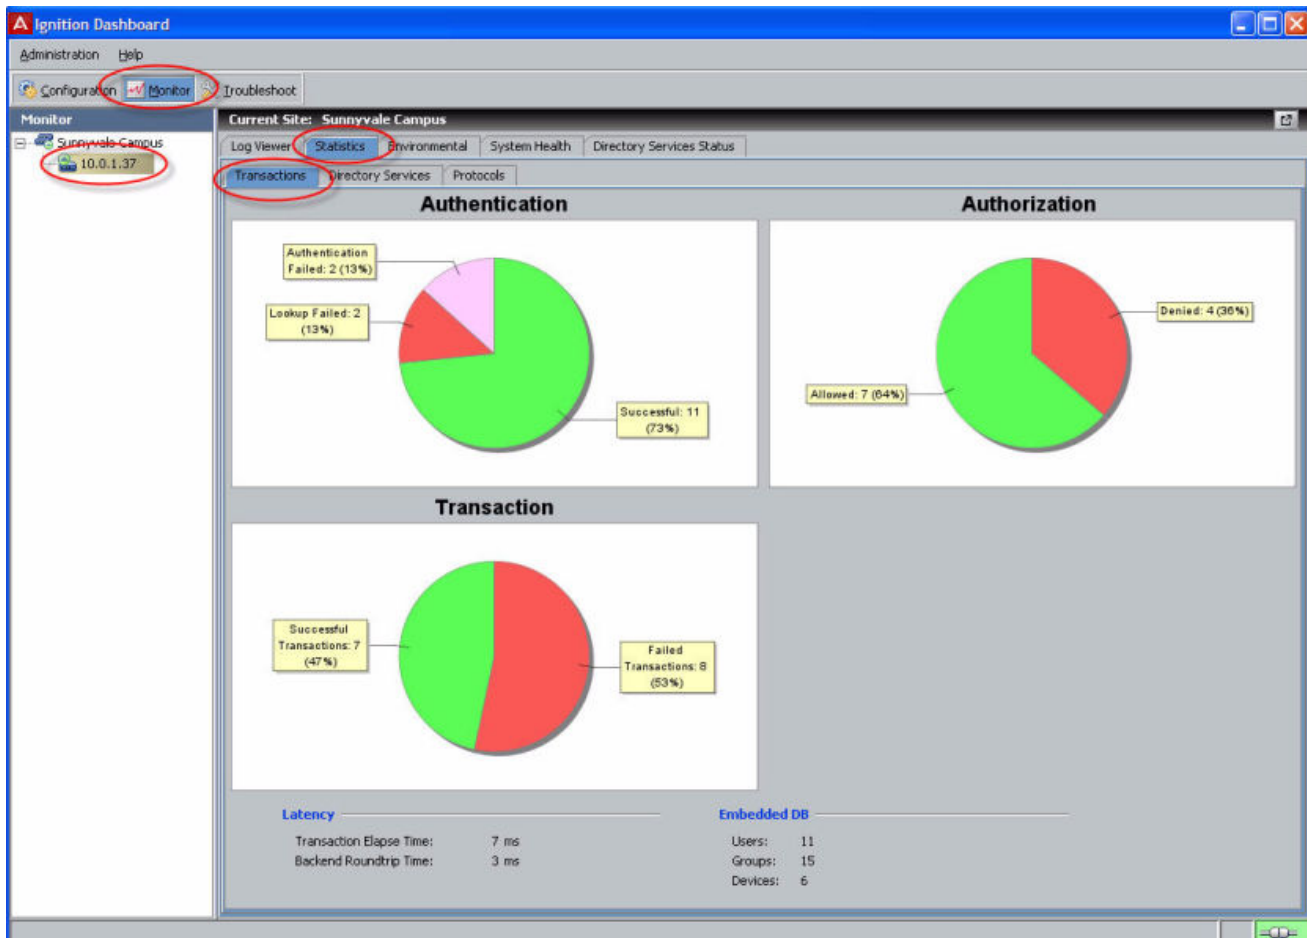


The Statistics tab lets you monitor the operation of the Ignition Server. Statistics are shown for the selected node only; click another node in the hierarchy tree to see its statistics. The display is refreshed every 5 seconds, and counters display the count since the last reboot of the Ignition Server.

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 pulldown 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.

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.
- **Embedded DB:** showing counts of users and groups in Ignition Server's local data store.


Viewing the Transactions Tab

Procedure

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

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

Procedure

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

Clicking the **Protocols** tab displays a set of sub-tabs, one per protocol, each offering statistics on Ignition Server's communications using the selected protocol.


For example, the **RADIUS** sub-tab of the **Protocols** tab provides detailed information about the number of RADIUS packets Ignition Server has received and processed (received, sent, accepted, rejected, and so on).

Viewing the Protocols Tab

Procedure

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 **Protocols** tab.

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.

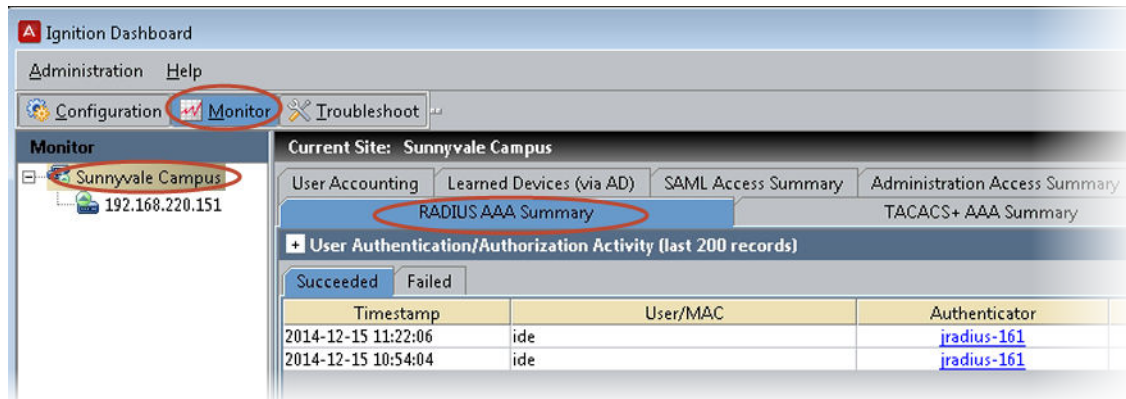
Procedure

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 186.

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 451.

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 lookup 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 a 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 more information, see [Viewing and managing logs](#) on page 451.

Viewing the AAA Summary Tabs

Follow this procedure to open the AAA Summary tabs.

Procedure

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 451.



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 display the most recent set of login attempts. Establish the maximum number of records to be displayed by configuring the value in the **Preferences** window. See [Setting viewing preferences for the Monitor view](#) on page 52.

The limit on the total number of entries is enforced across all three tabs; a tab might be empty if the other tabs contain enough recent records to reach the limit you configured in the **Preferences** window.

User Accounting tab

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

The main table in the User Accounting window displays a set of RADIUS attributes for each active session.

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.

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.



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 454.

Refresh button

To load the latest session audit data in the User Accounting tab, click the **Refresh** button.

Viewing the User Accounting Tab

Procedure

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

The Learned Devices tab displays a list of devices that have authenticated to Ignition Server using 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 474.

The expiration date and time for each device’s authentication is displayed in the **Expires** column. Each authentication lasts for the device Time To Live (TTL) period configured in the Learned Device TTL window. See [Setting TTL for Windows Machine authentication](#) on page 325.

Use the **Back** and **Next** buttons to move through the list. To filter the list, see [Filtering the Learned Devices Tab](#) on page 474.

Viewing the Learned Devices Tab

Procedure

1. Click **Monitor** in the main Dashboard window.
2. Click your site name in the **Monitor** hierarchy tree.
3. Click the **Learned Devices** tab.
4. If you want to filter the contents of the tab, see [Filtering the Learned Devices Tab](#) on page 474.

Filtering the Learned Devices Tab

You can filter the list of Learned Devices by selecting the **Specify Criteria** check box and:

- typing a full or partial **MAC Address** to be matched.
- 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.

Revoking the session of a Machine-Authenticated Device

To revoke the session of a machine-authenticated device, perform one of the following actions.

- To revoke a specific device session, click on its row to select it and click **Delete**.
- 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 a System 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.

 **Important:**

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.

SAML Access Summary tab

Contents of the SAML Access Summary tab

The SAML Access Summary gives a consolidated picture of the various SAML requests (successful and failed) processed by the Ignition Server. A separate view for successful and failed requests makes a clear distinction between how the SAML requests are processed and presents the user with the end outcome of each SAML request processing.

Viewing the SAML Access Summary tab

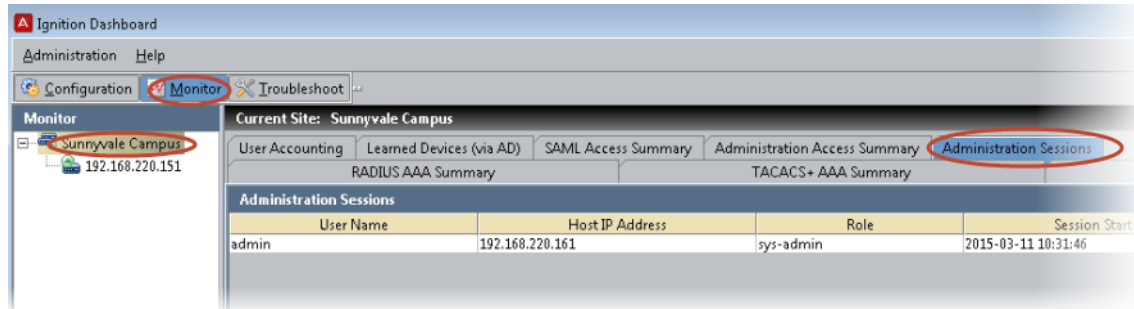
Procedure

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 **SAML Access Summary** tab.

Administration Sessions tab

The **Administration Sessions** tab shows the currently-active administrator sessions, including the following information:

- the user name
- the IP address from where they are logged in
- the administrator's role
- the session start and end times



Administration Access Summary tab

The **Administration Access Summary** tab shows 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 Administration Access Summary displays the following information.

- **Login Time:** the time when the administrator logged in
- **User Name:** User name of the administrator
- **Hostname:** The machine from where the dashboard was launched
- **Directory:** the name of the directory service that authenticated the user
- **Role:** Role of the administrator (sys-admin, config admin, monitor admin or troubleshoot admin)
- **Policy:** Policy that authenticated and authorized the user

* Note:

You can adjust the width of each column as necessary.

Failed tab

The **Failed** tab of the Administration Access Summary displays the following information.

- **Login Time:** the time when the administrator attempted to log in

- **User Name:** the User name of the administrator
- **Hostname:** the machine from where the dashboard was launched
- **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 lookup failed, this column shows the name of the last-searched directory service in your directory set.
- **Reason for Failure:** this column displays a 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.

Appendix H: Troubleshooting


This appendix explains how to generate trouble tickets and lists solutions for common errors that can occur when configuring Avaya Identity Engines Ignition Server.

Generating a trouble ticket


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.

Procedure

1. In Dashboard's Configuration hierarchy tree, right-click the name of your site and select **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 and click **Save**.
4. Click **Create Ticket**.



5. Ignition Server displays a progress bar. On completion, the following message is displayed.



- Contact technical support for instructions on uploading the file to Avaya. See [Support](#) on page 21.

Troubleshooting common problems

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.



If you receive this warning, you must replace the Ignition Engines 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 88.

Concurrent System Administrator Sessions not Allowed

The Ignition Server permits only one System Administrator session at a time. If a System Administrator is logged in, that administrator must log off before you can log in as System Administrator.

Problem: Connecting Dashboard to Ignition Server Fails

1. Details

When you attempt to connect to Dashboard, the connection attempt fails with the error “The Ignition Server is incompatible with the UI” displayed. This occurs if the firmware version on the Ignition Server is not supported by the current version of Ignition Dashboard.

1. Solution

Use the Ignition Server console 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.

2. Details

When a non system-admin user attempts to connect to Dashboard, the connection attempt fails with the error “Access not Allowed” displayed. This occurs if the Role Based Access Control (RBAC) feature is not enabled on the Ignition Server.

2. Solution

A System Administrator can log in to the Dashboard and go to **Site Configuration > Administration > Dashboard Hosts** and see if the global-default-mgmt host is enabled.

3. Details

When a non-System Administrator user attempts to connect to Dashboard, the connection attempt fails with the error “Role is not defined” displayed. This occurs if the user is not assigned any role as part of the policy evaluation.

3. Solution

A System Administrator can log in to the Dashboard and go to **Site Configuration > Administration > Admin Access Policies** and see if the policy associated with the global-default-mgmt host has the role definition clearly specified.

4. Details

When a user attempts to connect to Dashboard using the Site Group option, the connection attempt fails with the error “Unknown” displayed. This occurs if the connected node is unreachable.

4. Solution

Make sure that the Node to which the connection is attempted is up and reachable from the Dashboard. If the Node is up and reachable and this error still displays, the System Administrator can log in from the Dashboard to the Node and get more information from the Access Logs.

5. Details

When a user attempts to connect to Dashboard using the Site Group option, the connection attempt fails with the error “System Admin already logged in” displayed. This occurs if a System Administrator is already logged in to the Node to which you want to connect.

*** Note:**

Only a user with System Administrator credentials can log in using the Site Group option.

5. Solution

Make sure that the other System Administrator logs off from the Node to which the connection is attempted and retry this operation.

6. Details

When a user attempts to connect to Dashboard using the Site Group option, the connection attempt fails with the error “Non System Admin Credentials provided” displayed. This occurs if the credentials provided to connect to the Node are not those of a System Administrator.

6. Solution

Make sure that when creating a Site group, the credentials of System Administrator are provided for each Node.

7. Details

When a user attempts to connect to Dashboard, the connection attempt fails with the error `Session already exists (sys-admin)` displayed. This occurs if a System Administrator is already logged in to the Node to which you want to connect using another instance of Dashboard.

*** Note:**

Only one System Administrator session at a time is allowed on a node.

7. Solution

Make sure that there is no Dashboard instance with a valid System Administrator logged into the node to which you want to connect.

If for some reason, you are receiving the error `Session already exists (sys-admin)` even though there is no Dashboard instance with a valid System Administrator logged in, the session can be cleaned up on the server using the following commands on CLI.

```
show session
session delete <id>
session delete all
```

8. Details

When a user attempts to connect to Dashboard, the connection attempt fails with the error `Session already exists (cfg-admin)` displayed. This occurs if a Configuration Administrator is already logged in to the Node to which you want to connect using another instance of Dashboard.

*** Note:**

Only one Configuration Administrator session at a time is allowed on a node.

8. Solution

Make sure that there is no Dashboard instance with a valid Configuration Administrator logged into the node to which you want to connect.

If for some reason, you are receiving the error `Session already exists (cfg-admin)` even though there is no Dashboard instance with a valid Configuration Administrator logged in, the session can be cleaned up on the server using the following commands on CLI.

```
show session
session delete <id>
session delete all
```

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: [Configuring Ignition Server's RADIUS service](#) on page 57
 - for TACACS+: [Turning on the Ignition Server TACACS+ service](#) on page 329
2. Check the **Log Viewer: Security** tab to see if Ignition Server dropped the request. See [Security Log](#) on page 462. 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 104.
3. Check the **System Health** tab to make sure Ignition Server's **RADIUS Engine** is running. See [System Health tab](#) on page 468. If the RADIUS Engine is *not* running, contact Avaya customer support for help.

Problem: Authorization policy stops working unexpectedly

Possible Cause

Underlying but required 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.

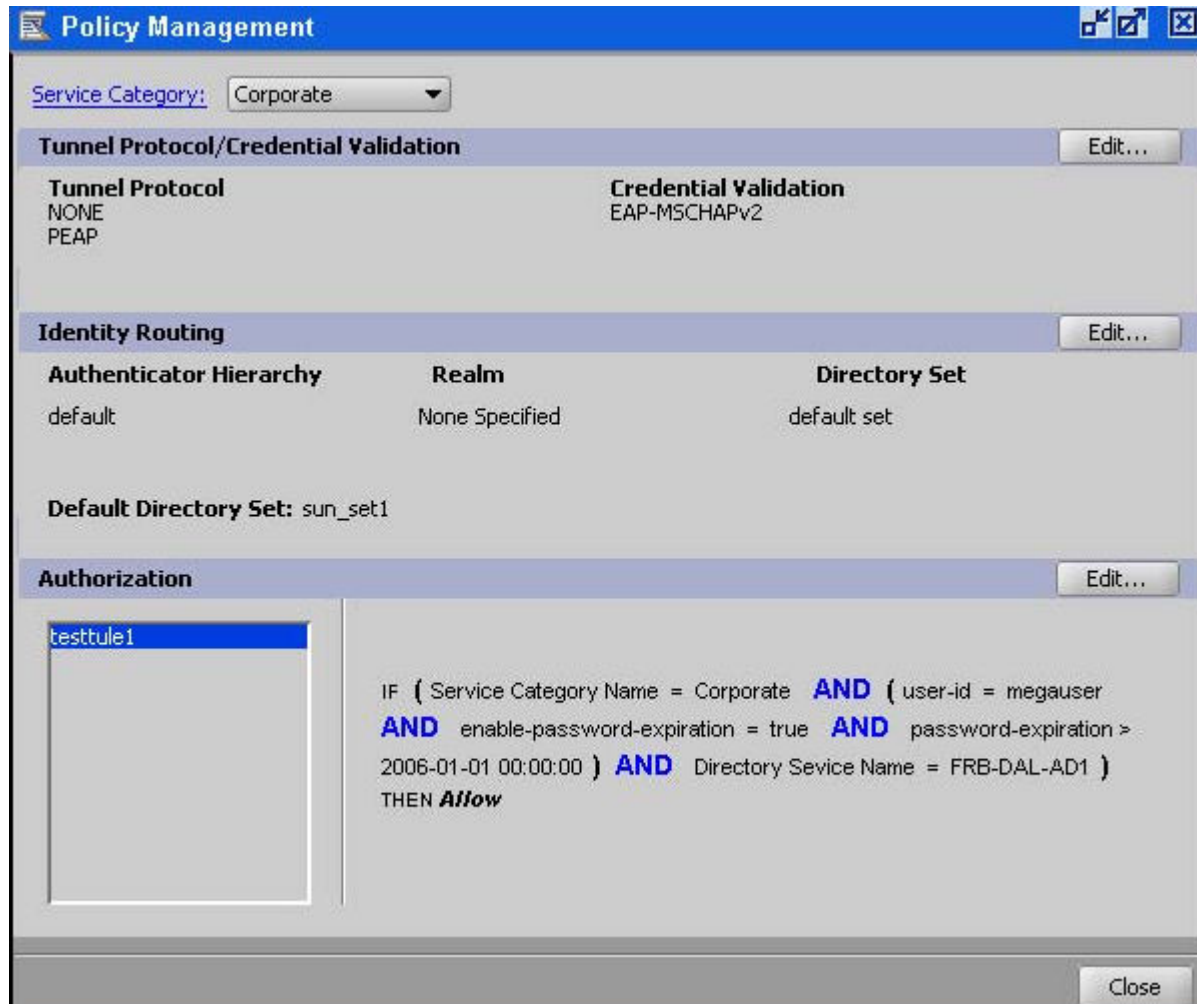
Each authorization policy can use data associated with the following Ignition Server data elements:

- authenticators
- authenticator bundles
- authenticator hierarchy (containers)
- access policies
- directory services within directory sets

- virtual groups
- virtual attributes

Example

The following figure 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.



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 following sections explain common failures that can occur in Active Directory environments. For more general authentication-related troubleshooting, see [Troubleshooting user lookup and authentication](#) on page 186.

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 calls to the AD server Netlogon service fail.

Solution

Edit your firewall settings as explained in [Preparing to connect to Active Directory](#) on page 153.

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 force Ignition Server to rejoin the domain as follows.

1. In Ignition Dashboard, click **Monitor** 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](#) on page 392, a bad network route might be the cause.

Solution

Fix this as follows.

1. Cancel the HA Configuration Wizard.
2. Repair the network connection between the HA port on your first Ignition Server and the HA port on your second Ignition Server.
3. From the current session of Dashboard, ping the HA port of the second Ignition Server.

Click **Troubleshoot** at the top of the Dashboard window and 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, configure 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.

4. Launch a new session of Dashboard and log on to the second Ignition Server. Leave the existing Dashboard connected to the first Ignition Server running.
5. Perform another ping test.

Click **Troubleshoot** at the top of the Dashboard window and 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. From the existing Dashboard session connected to the first Ignition Server, use the HA configuration wizard and repeat the steps to create the HA pair.

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,**" follow this procedure. In this procedure, the Ignition Servers are referred to 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 **Troubleshoot** 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, configure 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 Servers. 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 **Troubleshoot** 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, configure 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 Servers. 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 [Generating a trouble ticket](#) on page 478.

Problem: Errors occur during Directory Service Set-Up

Possible Cause

Ignition Server failed to parse your directory schema.

Details

When you click **Test Connections** 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 in 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 are *not* able to map virtual attributes to the directory, but you *can* authenticate against the directory and map virtual groups to it. See [Troubleshooting user lookup and authentication](#) on page 186.

Solution

- If this error occurs and you *do not* plan to use virtual attributes, then 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 [Support](#) on page 21.

Problem: Unable to Map Virtual Attributes

See [Problem: Errors occur during Directory Service Set-Up](#) on page 486.

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

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.
authentication	The process of verifying a user's (or device's) credentials to confirm their identity.
authentication server	<ol style="list-style-type: none">1. (<i>Avaya Identity Engines Ignition Server usage</i>) 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 <i>authentication decision</i> by evaluating the user credentials, and the directory server provides user attributes and group associations that form the basis for <i>authorization decision</i> 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. (<i>general usage</i>) 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	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.

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 on page 492.)
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 / EAPMSCHAPv2 authentication.” Stands for, “Extensible Authentication Protocol, Microsoft Challenge Handshake Authentication Protocol Version 2.”

groups	Labeled collections of users.
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 <i>Avaya Identity Engines Ignition Guest Manager Configuration</i> , NN47280-501 for details.
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 paired server high availability deployment (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 238.
site	In Ignition Server terminology, a <i>site</i> 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.