Dell Networking W-Series ArubaOS 6.4.x



User Guide

Copyright Information

© 2014 Aruba Networks, Inc. Aruba Networks trademarks include  , Aruba Networks®, Aruba Wireless Networks®, the registered Aruba the Mobile Edge Company logo, and Aruba Mobility Management System®. Dell™, the DELL™ logo, and PowerConnect™ are trademarks of Dell Inc.

All rights reserved. Specifications in this manual are subject to change without notice.

Originated in the USA. All other trademarks are the property of their respective owners.

Open Source Code

Certain Aruba products include Open Source software code developed by third parties, including software code subject to the GNU General Public License (GPL), GNU Lesser General Public License (LGPL), or other Open Source Licenses. Includes software from Litech Systems Design. The IF-MAP client library copyright 2011 Infoblox, Inc. All rights reserved. This product includes software developed by Lars Fenneberg, et al. The Open Source code used can be found at this site:

arubanetworks.com/open_source

Legal Notice

The use of Aruba Networks, Inc. switching platforms and software, by all individuals or corporations, to terminate other vendors' VPN client devices constitutes complete acceptance of liability by that individual or corporation for this action and indemnifies, in full, Aruba Networks, Inc. from any and all legal actions that might be taken against it with respect to infringement of copyright on behalf of those vendors.

Contents 3

About this Guide 85

What's New In ArubaOS 6.4.x 85

What's New In ArubaOS 6.4.0.0 89

Fundamentals 91

WebUI 91

CLI 91

Related Documents 92

Conventions 92

Contacting Dell 93

The Basic User-Centric Networks 94

Understanding Basic Deployment and Configuration Tasks 94

Deployment Scenario #1: Controller and APs on Same Subnet 94

Deployment Scenario #2: APs All on One Subnet Different from Controller Subnet 95

Deployment Scenario #3: APs on Multiple Different Subnets from Controllers 96

Configuring the Controller 97

Running Initial Setup 97

Connecting to the Controller after Initial Setup 98

Dell W-7200 Series Controller 98

New Port Numbering Scheme 98

Individual Port Behavior 99

Using the LCD Screen 99

Using the LCD and USB Drive 100

Upgrading an Image 101

Uploading a Pre-saved Configuration 101

Disabling LCD Menu Functions 101
Configuring a VLAN to Connect to the Network 102
Creating, Updating, and Viewing VLANs and Associated IDs 102
Creating, Updating, and Deleting VLAN Pools 103
Assigning and Configuring the Trunk Port 103
In the WebUI 103
In the CLI 103
Configuring the Default Gateway 103
In the WebUI 103
In the CLI 104
Configuring the Loopback IP Address for the Controller 104
In the WebUI 104
In the CLI 105
Configuring the System Clock 105
Installing Licenses 105
Connecting the Controller to the Network 105
Enabling Wireless Connectivity 105
Configuring Your User-Centric Network 106
Control Plane Security 107
Control Plane Security Overview 107
Configuring Control Plane Security 108
In the WebUI 108
In the CLI 110
Managing AP Whitelists 110
Adding APs to the Campus and Remote AP Whitelists 110
Viewing Whitelist Status 112
Modifying an AP in the Campus AP Whitelist 114
Revoking an AP via the Campus AP Whitelist 115

Deleting an AP Entry from the Campus AP Whitelist 115
Purging the Campus AP Whitelist 115
OffLoading a Controller RAP Whitelist to ClearPass Policy Manager 116
In the WebUI 116
In the CLI 117
Managing Whitelists on Master and Local Controllers 117
Campus AP Whitelist Synchronization 118
Viewing and Managing the Master or Local Controller Whitelists 118
Viewing the Master or Local Controller Whitelist 118
Deleting an Entry from the Master or Local Controller Whitelist 119
Purging the Master or Local Controller Whitelist 120
Working in Environments with Multiple Master Controllers 120
Configuring Networks with a Backup Master Controller 120
Configuring Networks with Clusters of Master Controllers 121
Creating a Cluster Root 121
Creating a Cluster Member 122
Viewing Controller Cluster Settings 122
Replacing a Controller on a Multi-Controller Network 123
Replacing Controllers in a Single Master Network 123
Replacing a Local Controller 123
Replacing a Master Controller with No Backup 124
Replacing a Redundant Master Controller 124
Replacing Controllers in a Multi-Master Network 125
Replacing a Local Controller in a Multi-Master Network 125
Replacing a Cluster Member Controller with no Backup 125
Replacing a Redundant Cluster Member Controller 126
Replacing a Cluster Root Controller with no Backup Controller 126
Replacing a Redundant Cluster Root Controller 126
Configuring Control Plane Security after Upgrading 127

Troubleshooting Control Plane Security 128
Identifying Certificate Problems 128
Verifying Certificates 128
Disabling Control Plane Security 128
Verifying Whitelist Synchronization 129
Rogue APs 129
Software Licenses 130
Understanding License Terminology 130
Working with Licenses 131
Centralized Licensing in a Multi-Controller Network 132
Primary and Backup Licensing Servers 133
Communication between the License Server and License Clients 133
Supported Topologies 135
Unsupported Topologies 136
Adding and Deleting Licenses 137
Replacing a Controller 137
Failover Behaviors 137
Client is Unreachable 138
Server is Unreachable 138
Configuring Centralized Licensing 138
Pre-configuration Setup in an All-Master Deployment 138
Preconfiguration Setup in a Master/Local Topology 139
Enabling Centralized Licensing 139
Using the WebUI 139
Using the CLI 139
Monitoring and Managing Centralized Licenses 140
License server Table 140
License Client Table 140
License Client(s) Usage Table 141

Aggregate License Table 141
License Heartbeat Table 142
Using Licenses 142
Understanding License Interaction 143
License Installation Best Practices and Exceptions 144
Installing a License 144
Enabling a new license on your controller 145
Requesting a Software License in Email 145
Locating the System Serial Number 145
Obtaining a Software License Key 145
Creating a Software License Key 146
Applying the Software License Key in the WebUI 146
Applying the Software License Key in the License Wizard 146
Deleting a License 146
Moving Licenses 147
Resetting the Controller 147
Network Configuration Parameters 148
Configuring VLANs 148
Creating and Updating VLANs 148
In the WebUI 148
In the CLI 149
Creating Bulk VLANs In the WebUI 149
In the CLI 149
Creating a VLAN Pool 149
Using the WebUI 149
Distinguishing Between Even and Hash Assignment Types 150
Updating a VLAN Pool 151
Deleting a VLAN Pool 151

Creating a VLAN Pool Using the CLI 151
Viewing and Adding VLAN IDs Using the CLI 151
Role Derivation for Named VLAN Pools 152
In the CLI 152
In the WebUI 152
Creating a Named VLAN not in a Pool 153
In the WebUI 153
In the CLI 153
Adding a Bandwidth Contract to the VLAN 154
Optimizing VLAN Broadcast and Multicast Traffic 154
Using the CLI 154
Using the WebUI 155
Configuring Ports 155
Classifying Traffic as Trusted or Untrusted 155
About Trusted and Untrusted Physical Ports 155
About Trusted and Untrusted VLANs 155
Configuring Trusted/Untrusted Ports and VLANs 156
In the WebUI 156
In the CLI 156
Configuring Trusted and Untrusted Ports and VLANs in Trunk Mode 157
In the WebUI 157
In the CLI 157
Understanding VLAN Assignments 157
VLAN Derivation Priorities for VLAN types 158
How a VLAN Obtains an IP Address 159
Assigning a Static Address to a VLAN 159
In the WebUI 159
In the CLI 159
Configuring a VLAN to Receive a Dynamic Address 159

Configuring Multiple Wired Uplink Interfaces (Active-Standby) 159
Enabling the DHCP Client 160
In the WebUI 160
In the CLI 160
Enabling the PPPoE Client 161
In the WebUI 161
In the CLI 161
Default Gateway from DHCP/PPPoE 161
In the WebUI 161
In the CLI 161
Configuring DNS/WINS Server from DHPC/PPPoE 161
In the WebUI 161
In the CLI 162
Configuring Source NAT to Dynamic VLAN Address 162
In the WebUI 162
In the CLI 162
Configuring Source NAT for VLAN Interfaces 163
Example Configuration 163
In the WebUI 163
In the CLI 163
Inter-VLAN Routing 164
Using the WebUI to restrict VLAN routing 164
Using the CLI 164
Configuring Static Routes 165
In the WebUI 165
In the CLI 165
Configuring the Loopback IP Address 165
In the WebUI 165
In the CLI 166

Configuring the Controller IP Address 166
Using the CLI 167
Configuring GRE Tunnels 167
Important Points to Remember 167
Limitations 167
Creating a Tunnel Interface 167
In the WebUI 168
In the CLI 168
Directing Traffic into the Tunnel 169
Static Routes 169
Firewall Policy 169
In the WebUI 169
In the CLI 169
Tunnel Keepalives 169
In the WebUI 169
In the CLI 170
Configuring GRE Tunnel Group 170
Creating a Tunnel Group 170
In the WebUI 170
In the CLI 170
Jumbo Frame Support 171
Limitations for Jumbo Frame Support 172
Configuring Jumbo Frame Support 172
Using the WebUI 172
Using the CLI 172
Viewing the Jumbo Frame Support Status 173
IPv6 Support 175
Understanding IPv6 Notation 175
Understanding IPv6 Topology 175

Enabling IPv6 176
Enabling IPv6 Support for Controller and APs 176
Configuring IPv6 Addresses 178
In the WebUI 178
To Configure Link Local Address 178
To Configure Global Unicast Address 179
To Configure Loopback Interface Address 179
In the CLI 179
Configuring IPv6 Static Neighbors 179
In the WebUI 179
In the CLI 180
Configuring IPv6 Default Gateway and Static IPv6 Routes 180
In the WebUI 180
To Configure IPv6 Default Gateway 180
To Configure Static IPv6 Routes 180
In the CLI 180
Managing Controller IP Addresses 180
In the WebUI 180
In the CLI 180
Configuring Multicast Listener Discovery (MLD) 181
In the WebUI 181
To Modify IPv6 MLD Parameters 181
In the CLI 181
Dynamic Multicast Optimization 182
In the WebUI 182
Using the WEBUI 182
In the CLI 182
Limitations 182
Debugging an IPv6 Controller 183

In the WebUI 183
In the CLI 183
Provisioning an IPv6 AP 183
In the WebUI 183
In the CLI 184
Enhancements to IPv6 Support on AP 184
Filtering an IPv6 Extension Header (EH) 184
Configuring a Captive Portal over IPv6 184
Working with IPv6 Router Advertisements (RAs) 184
Configuring an IPv6 RA on a VLAN 185
Using WebUI 186
Using CLI 186
Configuring Optional Parameters for RAs 186
In the WebUI 187
In the CLI 187
RADIUS Over IPv6 188
In the CLI 188
In the WebUI 189
TACACS Over IPv6 189
In the CLI 189
In the WebUI 190
DHCPv6 Server 190
Points to Remember 190
DHCP Lease Limit 190
Configuring DHCPv6 Server 191
In the WebUI 191
In the CLI 192
Understanding ArubaOS Supported Network Configuration for IPv6 Clients 193
Supported Network Configuration 193

Understanding the Network Connection Sequence for Windows IPv6 Clients 193
Understanding ArubaOS Authentication and Firewall Features that Support IPv6 193
Understanding Authentication 194
Working with Firewall Features 194
Understanding Firewall Policies 196
Creating an IPv6 Firewall Policy 197
Assigning an IPv6 Policy to a User Role 198
Understanding DHCPv6 Passthrough/Relay 198
Managing IPv6 User Addresses 198
Viewing or Deleting User Entries 198
Understanding User Roles 199
Viewing Datapath Statistics for IPv6 Sessions 199
Understanding IPv6 Exceptions and Best Practices 199
Link Aggregation Control Protocol (LACP) 201
Understanding LACP Best Practices and Exceptions 201
Configuring LACP 201
In the CLI 202
In the WebUI 203
LACP Sample Configuration 203
OSPFv2 205
Understanding OSPF Deployment Best Practices and Exceptions 205
Understanding OSPFv2 by Example using a WLAN Scenario 206
WLAN Topology 206
WLAN Routing Table 206
Understanding OSPFv2 by Example using a Branch Office Scenario 207
Branch Office Topology 207
Branch Office Routing Table 208
Configuring OSPF 208

Exporting VPN Client Addresses to OSPF 210
In the WebUI 210
In the CLI 210
Sample Topology and Configuration 210
Remote Branch 1 211
Remote Branch 2 212
W-3200 Central Office Controller—Active 213
W-3200 Central Office Controller—Backup 214
Topology 216
Observation 216
Configuring W-3600-UP Controller 216
Configuring W-3600-DOWN Controller 218
Viewing the Status of Instant AP VPN 219
RAPNG AP-1 219
RAPNG AP-3 220
Tunneled Nodes 222
Understanding Tunneled Node Configuration 222
Configuring a Wired Tunneled Node Client 223
Configuring an Access Port as a Tunneled Node Port 224
Configuring a Trunk Port as a Tunneled Node Port 224
Show commands 224
Authentication Servers 225
Understanding Authentication Server Best Practices and Exceptions 225
Understanding Servers and Server Groups 225
Configuring Authentication Servers 226
Configuring a RADIUS Server 226
Using the WebUI 226
Using the CLI 227

RADIUS Server VSAs 228
RADIUS Server Authentication Codes 231
RADIUS Server Fully Qualified Domain Names 231
DNS Query Intervals 231
Using the WebUI 232
Using the CLI 232
Configuring an RFC-3576 RADIUS Server 232
Using the WebUI 232
Using the CLI 232
Configuring an LDAP Server 232
Using the WebUI 233
Using the CLI 234
Configuring a TACACS+ Server 234
Using the WebUI 234
Using the CLI 234
Configuring a Windows Server 235
Using the WebUI 235
Using the CLI 235
Managing the Internal Database 235
Configuring the Internal Database 235
Using the WebUI 236
Using the CLI 236
Managing Internal Database Files 236
Exporting Files in the WebUI 237
Importing Files in the WebUI 237
Exporting and Importing Files in the CLI 237
Working with Internal Database Utilities 237
Deleting All Users 237
Repairing the Internal Database 237

Configuring Server Groups 238
Configuring Server Groups 238
Using the WebUI 238
Using the CLI 238
Configuring Server List Order and Fail-Through 238
Using the WebUI 239
Using the CLI 239
Configuring Dynamic Server Selection 239
Using the WebUI 240
Using the CLI 241
Configuring Match FQDN Option 241
Using the WebUI 241
Using the CLI 241
Trimming Domain Information from Requests 241
Using the WebUI 242
Using the CLI 242
Configuring Server-Derivation Rules 242
Using the WebUI 243
Using the CLI 243
Configuring a Role Derivation Rule for the Internal Database 243
Using the WebUI 244
Using the CLI 244
Assigning Server Groups 244
User Authentication 244
Management Authentication 244
Using the WebUI 245
Using the CLI 245
Accounting 245
RADIUS Accounting 245

Using the WebUI 247
Using the CLI 247
RADIUS Accounting on Multiple Servers 247
Using the CLI: 247
Using the WebUI: 247
TACACS+ Accounting 247
Configuring Authentication Timers 248
Setting an Authentication Timer 248
Using the WebUI 248
Using the CLI 248
Authentication Server Load Balancing 249
Enabling Authentication Server Load Balancing Functionality 249
MAC-based Authentication 250
Configuring MAC-Based Authentication 250
Configuring the MAC Authentication Profile 250
Using the WebUI to configure a MAC authentication profile 251
Using the CLI to configure a MAC authentication profile 251
Configuring Clients 251
In the WebUI 251
In the CLI 252
802.1X Authentication 253
Understanding 802.1X Authentication 253
Supported EAP Types 254
Configuring Authentication with a RADIUS Server 254
Configuring Authentication Terminated on Controller 255
Configuring 802.1X Authentication 256
In the WebUI 256
In the CLI 261
Configuring and Using Certificates with AAA FastConnect 261

In the WebUI 262
In the CLI 262
Configuring User and Machine Authentication 262
Working with Role Assignment with Machine Authentication Enabled 262
Enabling 802.1x Supplicant Support on an AP 264
Prerequisites 264
Provisioning an AP as an 802.1X Supplicant 264
In the WebUI 264
In the CLI 265
Sample Configurations 265
Configuring Authentication with an 802.1X RADIUS Server 265
Configuring Roles and Policies 266
Creating the Student Role and Policy 266
In the WebUI 266
In the CLI 267
Creating the Faculty Role and Policy 267
Using the WebUI 267
In the CLI 268
Creating the Guest Role and Policy 268
In the WebUI 268
In the CLI 269
Creating Roles and Policies for Sysadmin and Computer 269
In the WebUI 269
In the CLI 270
Using the WebUI to create the computer role 270
Using the CLI to create the computer role 270
Creating an Alias for the Internal Network Using the CLI 270
Configuring the RADIUS Authentication Server 270
In the WebUI 270
In the CLI 271

Configuring 802.1X Authentication 271
In the WebUI 271
In the CLI 271
Configuring VLANs 272
In the WebUI 272
In the CLI 272
Configuring the WLANs 273
Configuring the Guest WLAN 273
In the WebUI 273
In the CLI 274
Configuring the Non-Guest WLANs 274
In the WebUI 274
In the CLI 275
Configuring Authentication with the Controller's Internal Database 275
Configuring the Internal Database 275
In the WebUI 276
In the CLI 276
Configuring a Server Rule Using the WebUI 276
Configuring a Server Rule Using the CLI 276
Configuring 802.1x Authentication 276
In the WebUI 276
In the CLI 277
Configuring VLANs 277
In the WebUI 277
In the CLI 278
Configuring WLANs 278
Configuring the Guest WLAN 278
In the WebUI 279
In the CLI 279

Configuring the Non-Guest WLANs 279
In the WebUI 280
In the CLI 280
Configuring Mixed Authentication Modes 281
In the CLI 281
Performing Advanced Configuration Options for 802.1X 281
Configuring Reauthentication with Unicast Key Rotation 281
In the WebUI 282
In the CLI 282
Application Single Sign-On Using L2 Authentication 282
Important Points to Remember 283
Enabling Application SSO 283
Configuring SSO IDP-Profiles 283
In the WebUI 283
In the CLI 284
Applying an SSO Profile to a User Role 284
In the WebUI 284
In the CLI 284
Selecting an IDP Certificate 284
In the WebUI 284
In the CLI 284
Stateful and WISPr Authentication 285
Working With Stateful Authentication 285
Working With WISPr Authentication 286
Understanding Stateful Authentication Best Practices 286
Configuring Stateful 802.1X Authentication 286
In the WebUI 286
In the CLI 287

Configuring Stateful NTLM Authentication 287
In the WebUI 287
In the CLI 288
Configuring Stateful Kerberos Authentication 288
In the WebUI 288
In the CLI 289
Configuring WISPr Authentication 289
In the WebUI 289
In the CLI 290
Certificate Revocation 292
Understanding OCSP and CRL 292
Configuring a Controller as OCSP and CRL Clients 292
Configuring an OCSP Controller as a Responder 293
Configuring the Controller as an OCSP Client 293
In the WebUI 293
In the CLI 295
Configuring the Controller as a CRL Client 295
In the WebUI 295
In the CLI 296
Configuring the Controller as an OCSP Responder 296
In the WebUI 296
In the CLI 297
Certificate Revocation Checking for SSH Pubkey Authentication 297
Configuring the SSH Pubkey User with RCP 297
In the WebUI 297
In the CLI 297
Displaying Revocation Checkpoint for the SSH Pubkey User 298
Configuring the SSH Pubkey User with RCP 298

In the WebUI 298
In the CLI 298
Removing the SSH Pubkey User 298
In the WebUI 298
In the CLI 298
Captive Portal Authentication 299
Understanding Captive Portal 299
Policy Enforcement Firewall Next Generation (PEFNG) License 299
Controller Server Certificate 300
Configuring Captive Portal in the Base Operating System 300
In the WebUI 301
In the CLI 302
Using Captive Portal with a PEFNG License 302
Configuring Captive Portal in the WebUI 303
Configuring Captive Portal in the CLI 305
Sample Authentication with Captive Portal 305
Creating a Guest User Role 305
Creating an Auth-guest User Role 306
Configuring Policies and Roles in the WebUI 306
Creating a Time Range 306
Creating Aliases 307
Creating an Auth-Guest-Access Policy 307
Creating an Block-Internal-Access Policy 308
Creating a Drop-and-Log Policy 309
Creating a Guest Role 309
Creating an Auth-Guest Role 310
Configuring Policies and Roles in the CLI 310
Defining a Time Range 310

Creating Aliases 310
Creating a Guest-Logon-Access Policy 311
Creating an Auth-Guest-Access Policy 311
Creating a Block-Internal-Access Policy 311
Creating a Drop-and-Log Policy 311
Creating a Guest-Logon Role 311
Creating an Auth-Guest Role 311
Configuring Guest VLANs 311
In the WebUI 312
In the CLI 312
Configuring Captive Portal Authentication Profiles 312
Modifying the Initial User Role 313
Configuring the AAA Profile 313
Configuring the WLAN 314
Managing User Accounts 314
Configuring Captive Portal Configuration Parameters 315
Enabling Optional Captive Portal Configurations 317
Uploading Captive Portal Pages by SSID Association 317
Changing the Protocol to HTTP 318
Configuring Redirection to a Proxy Server 319
Redirecting Clients on Different VLANs 320
Web Client Configuration with Proxy Script 320
Personalizing the Captive Portal Page 321
Creating and Installing an Internal Captive Portal 323
Creating a New Internal Web Page 323
Username Example 324
Password Example 324
FQDN Example 324
Basic HTML Example 324

Installing a New Captive Portal Page 325
Displaying Authentication Error Messages 325
Reverting to the Default Captive Portal 326
Configuring Localization 326
Customizing the Welcome Page 329
Customizing the Pop-Up box 330
Customizing the Logged Out Box 331
Creating Walled Garden Access 332
In the WebUI 332
In the CLI 332
Enabling Captive Portal Enhancements 333
Configuring the Redirect-URL 333
Configuring the Login URL 333
Defining Netdestination Descriptions 333
Configuring a Whitelist 334
Configuring the Netdestination for a Whitelist: 334
Associating a Whitelist to Captive Portal Profile 334
Applying a Captive Portal Profile to a User-Role 334
Verifying a Whitelist Configuration 334
Verifying a Captive Portal Profile Linked to a Whitelist 334
Verifying Dynamic ACLs for a Whitelist 335
Verifying DNS Resolved IP Addresses for Whitelisted URLs 336
Virtual Private Networks 337
Planning a VPN Configuration 337
Selecting an IKE protocol 338
Understanding Suite-B Encryption Licensing 338
Working with IKEv2 Clients 339
Understanding Supported VPN AAA Deployments 339

Working with Certificate Groups 340
Working with VPN Authentication Profiles 340
Configuring a Basic VPN for L2TP/IPsec in the WebUI 342
Defining Authentication Method and Server Addresses 342
Defining Address Pools 343
RADIUS Framed-IP-Address for VPN Clients 343
Enabling Source NAT 343
Selecting Certificates 343
Defining IKEv1 Shared Keys 344
Configuring IKE Policies 344
Setting the IPsec Dynamic Map 345
Finalizing WebUI changes 346
Configuring a Basic L2TP VPN in the CLI 346
Configuring a VPN for L2TP/IPsec with IKEv2 in the WebUI 346
Defining Authentication Method and Server Addresses 347
Defining Address Pools 347
Enabling Source NAT 347
Selecting Certificates 347
Configuring IKE Policies 348
Setting the IPsec Dynamic Map 349
In the WebUI 349
Finalizing WebUI changes 350
In the CLI 350
Configuring a VPN for Smart Card Clients 350
Working with Smart Card clients using IKEv2 351
Working with Smart Card Clients using IKEv1 351
Configuring a VPN for Clients with User Passwords 351
In the WebUI 352
In the CLI 352

Configuring Remote Access VPNs for XAuth 353
Configuring VPNs for XAuth Clients using Smart Cards 353
Configuring a VPN for XAuth Clients Using a Username and Password 354
Working with Remote Access VPNs for PPTP 355
In the WebUI 355
In the CLI 355
Working with Site-to-Site VPNs 355
Working with Third-Party Devices 356
Working with Site-to-Site VPNs with Dynamic IP Addresses 356
Understanding VPN Topologies 356
Configuring Site-to-Site VPNs 357
In the WebUI 357
In the CLI 358
Detecting Dead Peers 360
About Default IKE Policies 360
Working with VPN Dialer 361
Configuring VPN Dialer 361
In the WebUI 362
In the CLI 362
Assigning a Dialer to a User Role 362
In the WebUI 362
In the CLI 363
Roles and Policies 364
Configuring Firewall Policies 364
Working With Access Control Lists (ACLs) 365
Support for Desktop Virtualization Protocols 365
Creating a Firewall Policy 365
In the WebUI 367

In the CLI 368
Creating a Network Service Alias 368
In the WebUI 368
In the CLI 369
Creating an ACL White List 369
In the WebUI 369
Configuring the ACL White List in the WebUI 369
Configuring the White List Bandwidth Contract in the CLI 369
Configuring the ACL White List in the CLI 370
User Roles 370
In the WebUI 370
In the CLI 372
Assigning User Roles 372
Assigning User Roles in AAA Profiles 372
In the WebUI 372
In the CLI 372
Working with User-Derived VLANs 373
Understanding Device Identification 374
Configuring a User-derived VLAN in the WebUI 374
Configuring a User-derived Role or VLAN in the CLI 374
User-Derived Role Example 375
RADIUS Override of User-Derived Roles 376
Configuring a Default Role for Authentication Method 376
In the WebUI 376
In the CLI 376
Configuring a Server-Derived Role 376
Configuring a VSA-Derived Role 376
Understanding Global Firewall Parameters 377
Using AppRF 2.0 381

Enabling Deep Packet Inspection (DPI) 381
In the WebUI 381
In the CLI 381
Show Command Output 382
Configuring Policies for AppRF 2.0 382
How ACL Works with AppRF 382
Global Session ACL 382
Role Default Session ACL 382
Session ACL Examples 383
In the WebUI 384
In the CLI 384
Configuring Bandwidth Contracts for AppRF 2.0 384
Global Bandwidth Contract Configuration 384
In the CLI 384
Role-Specific Bandwidth Contracts 384
Using an Exclude List 384
In the WebUI 385
In the CLI 385
ClearPass Policy Manager Integration 386
Introduction 386
Important Points to Remember 386
Enabling Downloadable Role on a Controller 387
Using the WebUI 387
Using the CLI 387
Sample Configuration 387
CPPM Server Configuration 387
Adding a Device 387
Adding Enforcement Profile 388
Advanced Role Configuration Mode 389

Adding Enforcement Policy 390
Adding Services 392
Controller Configuration 393
Configuring CPPM Server on Controller 393
Configuring Server Group to include CPPM Server 394
Configuring 802.1X Profile 394
Configuring AAA Profile 394
Show AAA Profile 394
Virtual APs 395
Virtual AP Profiles 395
Configuring the Virtual AP Profile 396
Creating and Configuring a Profile 396
Selective Multicast Stream 400
Associating Other Profiles to the Virtual AP 400
Configuring a Virtual AP in the CLI 401
Associating a Virtual AP Profile to an AP or AP Group 402
In the WebUI 402
In the CLI 402
Excluding a Virtual AP Profile 402
In the WebUI 402
In the CLI 403
Virtual AP Configuration Workflow 403
Using the WebUI 403
Using the CLI 403
Radio Resource Management (802.11k) 404
Configuring the 802.11k Profile 404
In the WebUI 404
In the CLI 406

Configuring Radio Resource Management Information Elements 406
In the WebUI 406
In the CLI 407
Configuring Beacon Report Requests 408
In the WebUI 408
In the CLI 409
Configuring Traffic Stream Measurement Report Requests 409
In the WebUI 409
In the CLI 411
BSS Transition Management (802.11v) 411
Frame Types 411
802.11k and 802.11v clients 412
Fast BSS Transition (802.11r) 412
Important Points to Remember 412
Configuring Fast BSS Transition 412
In the WebUI 413
In the CLI 413
Troubleshooting Fast BSS Transition 414
SSID Profiles 414
SSID Profile Overview 415
Suite-B Cryptography 415
Wi-Fi Multimedia Protection 416
Management Frame Protection 416
Configuring the SSID Profile 416
In the WebUI 416
In the CLI 421
WLAN Authentication 421
Configuring an AAA Profile in the WebUI 421
Configuring an AAA Profile in the CLI 423

High-Throughput Virtual APs 424
Configuring the High-Throughput Radio Profile 424
In the WebUI 424
In the CLI 425
Configuring the High-Throughput SSID Profile 425
In the WebUI 425
In the CLI 428
Guest WLANs 428
Configuring a Guest VLAN 429
In the WebUI 429
In the CLI 429
Configuring a Guest Role 429
In the WebUI 429
In the CLI 430
Configuring a Guest Virtual AP 430
In the WebUI 430
In the CLI 431
Adaptive Radio Management (ARM) 432
ARM Feature Overviews 432
Configuring ARM Settings 432
ARM Troubleshooting 432
Understanding ARM 432
ARM Support for 802.11n 433
Monitoring Your Network with ARM 433
Maintaining Channel Quality 433
Configuring ARM Scanning 433
Understanding ARM Application Awareness 434
Client Match 434

ARM Coverage and Interference Metrics 435
Configuring ARM Profiles 435
Creating and Configuring a New ARM Profile 436
In the WebUI 436
In the CLI 442
Modifying an Existing Profile 443
Copying an Existing Profile 443
Deleting a Profile 443
Assigning an ARM Profile to an AP Group 443
In the WebUI 444
In the CLI 444
Using Multi-Band ARM for 802.11a/802.11g Traffic 444
Band Steering 445
Steering Modes 445
Enabling Band Steering 446
In the WebUI 446
In the CLI 446
Enabling Traffic Shaping 446
Enabling Traffic Shaping 447
In the WebUI 447
In the CLI 448
Enabling or Disabling the Hard Limit Parameter in Traffic Management Profile 448
Using the WebUI 448
Using the CLI 449
Spectrum Load Balancing 449
Reusing Channels to Control RX Sensitivity Tuning 449
Configuring Non-802.11 Noise Interference Immunity 450
Troubleshooting ARM 450
Too many APs on the Same Channel 451

Wireless Clients Report a Low Signal Level 451
Transmission Power Levels Change Too Often 451
APs Detect Errors but Do Not Change Channels 451
APs Don't Change Channels Due to Channel Noise 451
Wireless Intrusion Prevention 452
Working with the Reusable Wizard 452
Understanding Wizard Intrusion Detection 453
Understanding Wizard Intrusion Protection 454
Protecting Your Infrastructure 454
Protecting Your Clients 454
Monitoring the Dashboard 455
Detecting Rogue APs 456
Understanding Classification Terminology 456
Understanding Classification Methodology 457
Understanding Match Methods 457
Understanding Match Types 457
Understanding Suspected Rogue Confidence Level 458
Understanding AP Classification Rules 458
Understanding SSID specification 458
Understanding SNR specification 458
Understanding Discovered-AP-Count specification 459
Sample Rules 459
Understanding Rule Matching 459
Working with Intrusion Detection 459
Understanding Infrastructure Intrusion Detection 459
Detecting an 802.11n 40MHz Intolerance Setting 463
Detecting Active 802.11n Greenfield Mode 464
Detecting Ad hoc Networks 464

Detecting an Ad hoc Network Using a Valid SSID 464
Detecting an AP Flood Attack 464
Detecting AP Impersonation 464
Detecting AP Spoofing 464
Detecting Bad WEP Initialization 464
Detecting a Beacon Frame Spoofing Attack 464
Detecting a Client Flood Attack 464
Detecting a CTS Rate Anomaly 465
Detecting an RTS Rate Anomaly 465
Detecting Devices with an Invalid MAC OUI 465
Detecting an Invalid Address Combination 465
Detecting an Overflow EAPOL Key 465
Detecting Overflow IE Tags 465
Detecting a Malformed Frame-Assoc Request 465
Detecting Malformed Frame-Auth 465
Detecting a Malformed Frame-HT IE 466
Detecting a Malformed Frame-Large Duration 466
Detecting a Misconfigured AP 466
Detecting a Windows Bridge 466
Detecting a Wireless Bridge 466
Detecting Broadcast Deauthentication 466
Detecting Broadcast Disassociation 466
Detecting Netstumbler 466
Detecting Valid SSID Misuse 466
Detecting Wellenreiter 466
Understanding Client Intrusion Detection 467
Detecting a Block ACK DoS 469
Detecting a ChopChop Attack 469
Detecting a Disconnect Station Attack 469

Detecting an EAP Rate Anomaly 469
Detecting a FATA-Jack Attack Structure 469
Detecting a Hotspotter Attack 470
Detecting a Meiners Power Save DoS Attack 470
Detecting an Omerta Attack 470
Detecting Rate Anomalies 470
Detecting a TKIP Replay Attack 470
Detecting Unencrypted Valid Clients 470
Detecting a Valid Client Misassociation 470
Detecting an AirJack Attack 471
Detecting ASLEAP 471
Detecting a Null Probe Response 471
Configuring Intrusion Protection 471
Understanding Infrastructure Intrusion Protection 471
Protecting 40MHz 802.11 High Throughput Devices 473
Protecting 802.11n High Throughput Devices 473
Protecting Against Adhoc Networks 473
Protecting Against AP Impersonation 473
Protecting Against Misconfigured APs 473
Protecting Against Wireless Hosted Networks 473
Protecting SSIDs 474
Protecting Against Rogue Containment 474
Protecting Against Suspected Rogue Containment 474
Protection against Wired Rogue APs 474
Understanding Client Intrusion Protection 474
Protecting Valid Stations 474
Protecting Windows Bridge 475
Configuring the WLAN Management System (WMS) 475
In the WebUI 475

In the CLI 476
Configuring Local WMS Settings 476
Managing the WMS Database 476
Understanding Client Blacklisting 476
Methods of Blacklisting 477
Blacklisting Manually 477
Blacklisting by Authentication Failure 477
Enabling Attack Blacklisting 478
Setting Blacklist Duration 479
Removing a Client from Blacklisting 479
Working with WIP Advanced Features 479
Configuring TotalWatch 480
Understanding TotalWatch Channel Types and Qualifiers 480
Understanding TotalWatch Monitoring Features 480
Understanding TotalWatch Scanning Spectrum Features 481
Understanding TotalWatch Channel Dwell Time 481
Understanding TotalWatch Channel Visiting 481
Understanding TotalWatch Age out of Devices 482
Administering TotalWatch 482
Configuring Per Radio Settings 482
Configuring Per AP Setting 482
Licensing 483
Tarpit Shielding Overview 483
Configuring Tarpit Shielding 484
EnablingTarpit Shielding 484
Understanding Tarpit Shielding Licensing CLI Commands 484
Access Points (APs) 485
Basic Functions and Features 485

Naming and Grouping APs 486
Creating an AP group 487
In the WebUI 487
In the CLI 487
Assigning APs to an AP Group 488
In the WebUI 488
In the CLI 488
Understanding AP Configuration Profiles 488
AP Profiles 489
RF Management Profiles 489
Wireless LAN Profiles 490
Mesh Profiles 493
QoS Profiles 493
IDS Profiles 494
HA Group profiles 494
Other Profiles 494
Profile Hierarchy 494
Viewing Profile Errors 495
Before you Deploy an AP 495
Mesh AP Preconfiguration 495
Remote AP Preconfiguration 495
Enable Controller Discovery 495
Controller Discovery using DNS 496
Controller Discovery using ADP 496
Controller discovery using a DHCP Server 496
Enable DHCP to Provide APs with IP Addresses 497
In the WebUI 497
In the CLI 497
AP Provisioning Profiles 497

Defining an AP Provisioning Profile 497
Assigning Provisioning Profiles 499
Configuring Installed APs 500
Configuring an AP using the Provisioning Wizard 500
Configuring a AP using the WebUI 500
Configuring a Remote AP 501
Remote Authentication 501
RAP Configuration 502
Configuring a Mesh AP 502
Verifying the Configuration 502
Optional AP Configuration Settings 503
AP Installation Mode 503
Using the WebUI 503
Using the CLI 504
AP Name 504
Using the WebU 504
Using the CLI 504
Spanning Tree 504
Using the WebUI 505
Using the CLI 505
RTLS Server 505
In the WebUI 505
In the CLI 506
Important Points to Remember 506
AP Redundancy 506
Using the WebUI 506
Using the CLI 506
AP Maintenance Mode 507
Using the WebUI 507

Using the CLI 507
Energy Efficient Ethernet 507
Using the WebUI 507
Using the CLI 508
AP LEDs 508
Using the WebUI 508
Using the CLI 509
RF Management 509
802.11a and 802.11g RF Management Profiles 509
VHT Support on W-AP200 Series, W-AP210 Series, W-AP220 Series, and W-AP270 Series Access Points 510
Managing 802.11a/802.11g Profiles Using the WebUI 511
Creating or Editing a Profile 511
Assigning an 802.11a/802.11g Profile to an AP or AP Group 515
Assigning a High-throughput Profile 515
Assigning an ARM Profile 516
Deleting a Profile 517
Managing 802.11a/802.11g Profiles Using the CLI 517
Creating or Modifying a Profile 517
Viewing RF Management Settings 518
Assigning a 802.11a/802.11g Profile 518
Deleting a Profile 518
RF Optimization 518
Using the WebUI 518
Using the CLI 519
RF Event Configuration 519
Using the WebUI 519
Using the CLI 521
Optimizing APs Over Low-Speed Links 521

Configuring the Bootstrap Threshold 522
Prioritizing AP heartbeats 525
526
Configuring AP Channel Assignments 526
Using the WebUI 526
Using the CLI 527
Channel Switch Announcement (CSA) 527
Using the WebUI 528
Using the CLI 528
Automatic Channel and Transmit Power Selection 528
Managing AP Console Settings 528
Link Aggregation Support on W-AP220 Series and W-AP270 Series 529
Configuring LACP 530
Using the WebUI, in ArubaOS 6.4.2.x and later 530
Using the CLI, in ArubaOS 6.4.2.x and later 530
Using the WebUI in ArubaOS 6.3.1.x-6.4.1.x 530
Using the CLI in ArubaOS 6.3.1.x-6.4.1.x 531
Important Points to Remember 531
Troubleshooting Link Aggregation 531
Service Tag 531
In the WebUI 532
In the CLI 532
Secure Enterprise Mesh 533
Mesh Overview Information 533
Mesh Configuration Procedures 533
Understanding Mesh Access Points 533
Mesh Portals 534
Mesh Points 534

Mesh Clusters 535
Understanding Mesh Links 535
Link Metrics 536
Optimizing Links 537
Understanding Mesh Profiles 537
Mesh Cluster Profiles 537
Mesh Radio Profiles 538
RF Management (802.11a and 802.11g) Profiles 539
Adaptive Radio Management Profiles 539
High-Throughput Radio Profiles 540
Mesh High-Throughput SSID Profiles 540
Wired AP Profiles 540
Mesh Recovery Profiles 540
Understanding Remote Mesh Portals (RMPs) 541
Understanding the AP Boot Sequence 542
Booting the Mesh Portal 542
Booting the Mesh Point 543
Air Monitoring and Mesh 543
Mesh Deployment Solutions 543
Thin AP Services with Wireless Backhaul Deployment 543
Point-to-Point Deployment 544
Point-to-Multipoint Deployment 544
High-Availability Deployment 545
Mesh Deployment Planning 545
Pre-Deployment Considerations 546
Outdoor-Specific Deployment Considerations 546
Configuration Considerations 546
Post-Deployment Considerations 547
Dual-Port AP Considerations 547

Configuring Mesh Cluster Profiles 547
Managing Mesh Cluster Profiles in the WebUI 547
Creating a Profile 547
Associating a Mesh Cluster Profile to Mesh APs 549
Editing a Mesh Cluster Profile 549
Deleting a Mesh Cluster Profile 550
Managing Mesh Cluster Profiles in the CLI 550
Viewing Mesh Cluster Profile Settings 550
Associating Mesh Cluster Profiles 550
Excluding a Mesh Cluster Profile from a Mesh Node 551
Deleting a Mesh Cluster Profile 551
Creating and Editing Mesh Radio Profiles 551
Managing Mesh Radio Profiles in the WebUI 551
Creating or Editing a Mesh Radio Profile 551
Assigning a Mesh Radio Profile to a Mesh AP or AP Group 554
Managing Mesh Radio Profiles in the CLI 554
Creating or Modifying a Mesh Radio Profile 555
Assigning a Mesh Radio Profile to a Mesh AP or AP Group 555
Deleting Mesh Radio Profiles 555
Creating and Editing Mesh High-Throughput SSID Profiles 556
Managing Mesh High-Throughput SSID Profiles in the WebUI 556
Creating a Profile 556
Assigning a Profile to an AP Group 559
Editing a Profile 559
Deleting a Profile 560
Managing Mesh High-Throughput SSID Profiles in the CLI 560
Creating or Modifying a Profile 560
Assigning a Profile to an AP Group 560
Viewing High-throughput SSID Settings 561

Deleting a Profile 561
Configuring Ethernet Ports for Mesh 561
Configuring Bridging on the Ethernet Port 561
Configuring Ethernet Ports for Secure Jack Operation 562
In the WebUI 562
In the CLI 563
Extending the Life of a Mesh Network 563
In the WebUI 563
In the CLI 563
Provisioning Mesh Nodes 564
Provisioning Caveats 564
Provisioning Mesh Nodes 565
In the WebUI 565
In the CLI 565
Verifying Your Mesh Network 566
Verification Checklist 566
CLI Examples 566
Configuring Remote Mesh Portals (RMPs) 567
Creating a Remote Mesh Portal In the WebUI 568
Step 1: Provision the AP 568
Step 2: Define the Mesh Private VLAN in the Mesh Radio Profile 568
Step 3: Assign the Mesh Radio Profile to a Remote Mesh AP 568
Step 4: Assign an RF Management Profile to a Remote Mesh AP 569
Step 5: Assign a Mesh Cluster Profile 569
Step 6: Configuring a DHCP Pool 569
Step 7: Configuring the VLAN ID of the Virtual AP Profile 569
Provisioning a Remote Mesh Portal In the CLI 569

Increasing Network Uptime Through Redundancy and VRRP 570
High Availability 570
Pre-Deployment Information 570
Configuration Procedures 570
VRRP-Based Redundancy 570
High Availability Deployment Models 571
Active/Active Deployment Model 571
1:1 Active/Standby Deployment Model 571
N:1 Active/Standby Deployment Model 572
Master-Redundancy Deployment Model 572
AP Communication with Controllers 573
Client State Synchronization 573
Feature Guidelines and Limitations 574
High Availability Inter-Controller Heartbeats 574
High Availability Extended Controller Capacity 574
Feature Requirements 575
Standby Controller Capacity 575
AP Failover 576
Configuring High Availability 576
Pre-Deployment Information 576
Configuring High Availability 576
In the WebUI 576
In the CLI 577
Migrating from VRRP or Backup-LMS Redundancy 578
Configuring a Master Controller for Redundancy and High Availability: 578
Migrating from VRRP Redundancy 579
Migrating from Backup-LMS Redundancy 579
Configuring VRRP Redundancy 579

Before you Begin 580
Configuring the Local Controller for Redundancy 580
In the WebUI 580
In the CLI 581
Configuring the LMS IP 582
In the WebUI 582
In the CLI 582
Configuring the Master Controller for Redundancy 582
Configuring Database Synchronization 584
In the WebUI 584
In the CLI 584
Enabling Incremental Configuration Synchronization (CLI Only) 584
Configuring Master-Local Controller Redundancy 585
RSTP 587
Understanding RSTP Migration and Interoperability 587
Working with Rapid Convergence 587
Edge Port and Point-to-Point 589
Configuring RSTP 589
In the WebUI 589
In the CLI 590
Monitoring RSTP 590
Troubleshooting RSTP 591
PVST+ 593
Understanding PVST+ Interoperability and Best Practices 593
Enabling PVST+ in the CLI 593
Enabling PVST+ in the WebUI 594
Link Layer Discovery Protocol 595
Important Points to Remember 595

LLDP Overview 595
Default LLDP Configuration 596
Configuring LLDP 596
Monitoring LLDP Configuration 596
Display LLDP Interface 596
Display LLDP Interface <interface> 596
Display LLDP Neighbor 597
Display LLDP Neighbor Interface Detail 597
Display LLDP Statistics 598
Display LLDP Statistics Interface 598
IP Mobility 599
Understanding Dell Mobility Architecture 599
Configuring Mobility Domains 600
Configuring a Mobility Domain 601
Using the WebUI 601
Using the CLI 601
Joining a Mobility Domain 602
In the WebUI 602
In the CLI 602
Example Configuration 602
Configuring Mobility using the WebUI 603
Configuring Mobility using the CLI 603
Tracking Mobile Users 603
Mobile Client Roaming Status 604
Viewing Mobile Client Status using the WebUI 604
Viewing Mobile Client Status using the CLI 604
Viewing User Roaming Status using the CLI 604
Viewing specific client information using the CLI 605

Mobile Client Roaming Locations 605
In the WebUI 605
In the CLI 605
HA Discovery on Association 605
Setting up Mobility Association using the CLI 605
Configuring Advanced Mobility Functions 606
In the WebUI 606
In the CLI 608
Proxy Mobile IP 608
Revocations 608
IPv6 L3 Mobility 609
Multicast Mobility 609
Important Points to Remember 609
Example Configuration 611
Understanding Bridge Mode Mobility Deployments 615
Enabling Mobility Multicast 616
Working with Proxy IGMP and Proxy Remote Subscription 616
IGMPv3 Support 617
Configuring SSM Range 617
Using the CLI 617
Using the WebUI 617
Working with Inter Controller Mobility 618
Configuring Mobility Multicast 619
In the WebUI 619
In the CLI 619
Example 619
Palo Alto Networks Firewall Integration 620
Limitations 620
Preconfiguration on the PAN Firewall 620

User-ID Support 621
Device-Type Based Policy Support 621
Configuring PAN Firewall Integration 622
Creating PAN Profiles 622
Using the WebUI 623
Using the CLI 623
Activating a PAN Profile 623
Using the WebUI 624
Using the CLI 624
Enabling PAN Firewall Integration 624
Using the WebUI 624
Using the CLI 624
Enabling PAN Firewall Integration for VIA Clients 624
Using the WebUI 624
Using the CLI 624
Enabling PAN Firewall Integration for VPN Clients 624
Using the WebUI 625
Using the CLI 625
External Firewall Configuration 626
Understanding Firewall Port Configuration Among Dell Devices 626
Communication Between Controllers 626
Communication Between APs and the Controller 626
Communication Between Remote APs and the Controller 627
Enabling Network Access 627
Ports Used for Virtual Internet Access (VIA) 627
Configuring Ports to Allow Other Traffic Types 627
Remote Access Points 629
About Remote Access Points 629

Configuring the Secure Remote Access Point Service 631
Configure a Public IP Address for the Controller 631
Using the WebUI to create a DMZ address 631
Using CLI 631
Configure the NAT Device 632
Configure the VPN Server 632
Using the WebUI 632
Using CLI 632
CHAP Authentication Support over PPPoE 632
Using the WebUI to configure CHAP 632
Using the CLI to configure the CHAP 633
Configuring Certificate RAP 633
Using WebUI 633
Using CLI 633
Creating a Remote AP Whitelist 633
Configuring PSK RAP 634
Add the user to the internal database 634
Using WebUI 634
Using CLI 634
RAP Static Inner IP Address 635
Using the WebUI 635
Using the CLI 635
Provision the AP 635
Deploying a Branch Office/Home Office Solution 636
Provisioning the Branch Office AP 637
Configuring the Branch Office AP 637
Troubleshooting Remote AP 637
Local Debugging 638
Remote AP Summary 638

Multihoming on remote AP (RAP) 640
Seamless failover from backup link to primary link on RAP 640
Remote AP Connectivity 641
Remote AP Diagnostics 641
Enabling Remote AP Advanced Configuration Options 641
Understanding Remote AP Modes of Operation 642
Working in Fallback Mode 644
Backup Configuration Behavior for Wired Ports 645
Configuring Fallback Mode 645
Configuring the AAA Profile for Fallback Mode in the WebUI 645
Configuring the AAA Profile for Fallback Mode in the CLI 646
Configuring the Virtual AP Profile for Fallback Mode in the WebUI 646
Configuring the Virtual AP Profile for Fallback Mode in the CLI 647
Configuring the DHCP Server on the Remote AP 647
Using the WebUI 647
Using CLI 648
Configuring Advanced Backup Options 649
Configuring the Session ACL in the WebUI 649
Configuring the AAA Profile in the WebUI 650
Defining the Backup Configuration in the WebUI 650
Configuring the Session ACL in the CLI 651
Using the CLI to configure the AAA profile 651
Defining the Backup Configuration in the CLI 652
Specifying the DNS Controller Setting 652
In the WebUI 653
Backup Controller List 653
Configuring the LMS and backup LMS IP addresses in the WebUI 653
Configuring the LMS and backup LMS IP addresses in the CLI 654
Configuring Remote AP Failback 654

In the WebUI 654
In the CLI 654
Enabling RAP Local Network Access 654
In the WebUI 654
In the CLI 655
Configuring Remote AP Authorization Profiles 655
In the WebUI 655
Adding or Editing a Remote AP Authorization Profile 655
In the CLI 656
Working with Access Control Lists and Firewall Policies 656
Understanding Split Tunneling 656
Configuring Split Tunneling 657
Configuring the Session ACL Allowing Tunneling 657
Using the WebUI 657
Using the CLI 658
Configuring an ACL to Restrict Local Debug Homepage Access 659
In the WebUI 659
In the CLI 660
Configuring the AAA Profile for Tunneling 660
In the WebUI 660
Inthe CLI 661
Configuring the Virtual AP Profile 661
In the WebUI 661
In the CLI 661
Defining Corporate DNS Servers 662
In the WebUI 662
In the CLI 662
Understanding Bridge 662
Configuring Bridge 663

Configuring the Session ACL 663
Using the WebUI 663
Using the CLI 665
Configuring the AAA Profile for Bridge 665
In the WebUI 665
In the CLI 665
Configuring Virtual AP Profile 666
In the WebUI 666
In the CLI 666
Provisioning Wi-Fi Multimedia 667
Reserving Uplink Bandwidth 667
Understanding Bandwidth Reservation for Uplink Voice Traffic 667
Configuring Bandwidth Reservation 667
In the WebUI 667
In the CLI 668
Provisioning 4G USB Modems on Remote Access Points 668
4G USB Modem Provisioning Best Practices and Exceptions 668
Provisioning RAP for USB Modems 669
In the WebUI 669
In the CLI 669
RAP 3G/4G Backhaul Link Quality Monitoring 670
Provisioning RAPs at Home 670
Prerequisites 670
Provisioning RAP Using Zero Touch Provisioning 671
Provisioning the RAP using a Static IP Address 671
Provision the RAP on a PPPoE Connection 672
Using 3G/EVDO USB Modems 672
Configuring W-IAP3WN and W-IAP3WNP Access Points 674
Using the WebUI 674

Using the CLI 674
Converting an IAP to RAP or CAP 674
Converting IAP to RAP 675
Converting an IAP to CAP 675
Enabling Bandwidth Contract Support for RAPs 676
Configuring Bandwidth Contracts for RAP 676
Defining Bandwidth Contracts 676
Applying Contracts 676
Applying Contracts Per-Role 676
Applying Contracts Per-User 676
Verifying Contracts on AP 676
Verifying Contracts Applied to Users 677
Verifying Bandwidth Contracts During Data Transfer 678
Virtual Intranet Access 679
Understanding VIA Connection Manager 679
How it Works 679
Installing the VIA Connection Manager 680
On Microsoft Windows Computers 680
On Apple MacBooks 680
Upgrade Workflow 681
Minimal Upgrade 681
Complete Upgrade 681
VIA Compatibility Matrix 681
Configuring the VIA Controller 682
Before you Begin 682
Supported Authentication Mechanisms 682
Authentication mechanisms supported in VIA 1.x 682
Authentication mechanisms supported in VIA 2.x 682
Other authentication methods: 683

Suite B Cryptography Support 683
802.11 Suite-B 683
Configuring VIA Settings 683
Using the WebUI to Configure VIA 684
Enable VPN Server Module 684
Create VIA User Roles 684
Create VIA Authentication Profile 684
Create VIA Connection Profile 686
Configure VIA Web Authentication 690
Associate VIA Connection Profile to User Role 691
Configure VIA Client WLAN Profiles 691
Rebranding VIA and Downloading the Installer 694
Download VIA Installer and Version File 694
Customize VIA Logo 695
Customize the Landing Page for Web-based Login 695
Using the CLI to Configure VIA 695
Create VIA roles 695
Create VIA authentication profiles 695
Create VIA connection profiles 695
Configure VIA web authentication 696
Associate VIA connection profile to user role 696
Configure VIA client WLAN profiles 696
Customize VIA logo, landing page and downloading installer 696
Downloading VIA 696
Pre-requisites 696
Downloading VIA 697
Installing VIA 698
Using VIA 698
Connection Details Tab 698

Diagnostic Tab 699
Settings Tab 699
Troubleshooting 699
Spectrum Analysis 700
Understanding Spectrum Analysis 700
Spectrum Analysis Clients 703
Hybrid AP Channel Changes 704
Hybrid APs Using Mode-Aware ARM 704
Creating Spectrum Monitors and Hybrid APs 705
Converting APs to Hybrid APs 705
In the WebUI 705
In the CLI 705
Converting an Individual AP to a Spectrum Monitor 706
In the WebUI 706
In the CLI 706
Converting a Group of APs to Spectrum Monitors 706
In the WebUI 707
In the CLI 707
Connecting Spectrum Devices to the Spectrum Analysis Client 707
View Connected Spectrum Analysis Devices 708
Disconnecting a Spectrum Device 709
Configuring the Spectrum Analysis Dashboards 710
Selecting a Spectrum Monitor 710
Changing Graphs within a Spectrum View 711
Renaming a Spectrum Analysis Dashboard View 711
Saving a Dashboard View 712
Resizing an Individual Graph 713
Customizing Spectrum Analysis Graphs 713

Spectrum Analysis Graph Configuration Options 714
Active Devices 714
Active Devices Table 715
Active Devices Trend 718
Channel Metrics 719
Channel Metrics Trend 721
Channel Summary Table 723
Device Duty Cycle 724
Channel Utilization Trend 726
Devices vs Channel 727
FFT Duty Cycle 729
Interference Power 731
Quality Spectrogram 733
Real-Time FFT 734
Swept Spectrogram 736
Working with Non-Wi-Fi Interferers 739
Understanding the Spectrum Analysis Session Log 741
Viewing Spectrum Analysis Data 741
Recording Spectrum Analysis Data 742
Creating a Spectrum Analysis Record 742
Saving the Recording 743
Playing a Spectrum Analysis Recording 744
Playing a Recording in the Spectrum Dashboard 744
Playing a Recording Using the RFPlayback Tool 744
Troubleshooting Spectrum Analysis 745
Verifying Spectrum Monitors Support for One Client per Radio 745
Converting a Spectrum Monitor Back to an AP or Air Monitor 745
Troubleshooting Browser Issues 745
Loading a Spectrum View 746

Troubleshooting Issues with Adobe Flash Player 10.1 or Later 746
Understanding Spectrum Analysis Syslog Messages 746
Playing a Recording in the RFPlayback Tool 746
Dashboard Monitoring 747
Performance 747
Clients 747
APs 748
Using Dashboard Histograms 748
Usage 748
Security 749
Potential Issues 749
WLANs 750
Access Points 750
Clients 751
Firewall 752
In the WebUI 752
In the CLI 752
Element View 753
Details View 754
Element Tab 754
Element Summary View 754
Usage Breakdown 756
Aggregated Sessions 756
AppRF 758
All Traffic 758
Action Bar 759
Filters 759
Details 760

Block/Unblock, Throttle, and QoS Action Buttons 762
Block/Unblock 763
Throttle 765
QoS 765
Web Content Classification 766
Web Content Filters 769
WebCC Configuration in the WebUI 770
Block / Unblock: 770
Throttle 771
QoS 772
WebCC Configuration in the CLI 772
Enabling WebCC 772
New policy configuration 772
WebCC Bandwidth Contract Configuration 773
AirGroup 774
UCC 775
Chart View 776
Details View 776
Management Access 778
Configuring Certificate Authentication for WebUI Access 778
In the WebUI 778
In the CLI 779
Secure Shell (SSH) 779
Enabling Public Key Authentication 779
In the WebUI 780
In the CLI 780
Enabling RADIUS Server Authentication 780
Configuring RADIUS Server Username and Password Authentication 780
In the WebUI 780
In the CLI 781

Configuring RADIUS Server Authentication with VSA 781
Configuring RADIUS Server Authentication with Server Derivation Rule 781
In the WebUI 781
In the CLI 782
Configuring a set-value server-derivation rule 782
In the WebUI 782
In the CLI 783
Disabling Authentication of Local Management User Accounts 783
In the WebUI 783
In the CLI 783
Verifying the configuration 784
Resetting the Admin or Enable Password 784
Bypassing the Enable Password Prompt 785
Setting an Administrator Session Timeout 785
In the WebUI 785
In the CLI 785
Connecting to an AirWave Server 785
Custom Certificate Support for RAP 786
Suite-B Support for ECDSA Certificate 787
Setting the Default Server Certificate 787
In the CLI 787
Importing a Custom Certificate 787
In the WebUI 787
Generating a CSR 787
Uploading the Certificate 788
Implementing a Specific Management Password Policy 788
Defining a Management Password Policy 788
In the WebUI 788
Management Authentication Profile Parameters 789

Configuring AP Image Preload 790
Enable and Configure AP Image Preload 790
In the WebUI 791
In the CLI 791
View AP Preload Status 792
Configuring Centralized Image Upgrades 792
Configuring Centralized Image Upgrades 793
Using the WebUI 793
In the CLI 794
Viewing Controller Upgrade Statistics 794
Managing Certificates 795
About Digital Certificates 796
Obtaining a Server Certificate 796
In the WebUI 796
In the CLI 797
Obtaining a Client Certificate 797
Importing Certificates 798
In the WebUI 798
In the CLI 798
Viewing Certificate Information 798
Imported Certificate Locations 799
Checking CRLs 799
Certificate Expiration Alert 800
Chained Certificates on the RAP 800
Support for Certificates on USB Flash Drives 800
Marking the USB Device Connected as a Storage Device 800
RAP Configuration Requirements 800
Configuring SNMP 801
SNMP Parameters for the Controller 801

In the WebUI 802
In the CLI 802
Enabling Capacity Alerts 802
In the WebUI 803
In the CLI 803
Examples 803
Configuring Logging 804
In the WebUI 805
In the CLI 806
Enabling Guest Provisioning 806
Configuring the Guest Provisioning Page 806
In the WebUI 806
Configuring the Guest Fields 807
Configuring the Page Design 809
Configuring Email Messages 809
Configuring the SMTP Server and Port in the WebUI 810
Configuring an SMTP server and port in the CLI 810
Creating Email Messages in the WebUI 810
Configuring a Guest Provisioning User 811
In the WebUI 811
Username and Password Authentication Method 811
Static Authentication Method 811
Smart Card Authentication Method 812
In the CLI 812
Username and Password Method 812
Static Authentication Method 812
Smart Card Authentication Method 812
Customizing the Guest Access Pass 813
Creating Guest Accounts 813
Guest Provisioning User Tasks 814

Importing Multiple Guest Entries 815
Creating Multiple Guest Entries in a CSV File 815
Importing the CSV File into the Database 816
Printing Guest Account Information 819
Optional Configurations 820
Restricting one Captive Portal Session for each Guest 820
Using the CLI to restrict one Captive Portal session for each guest 820
Setting the Maximum Time for Guest Accounts 820
Using the WebUI to set the maximum time for guest accounts 821
Using the CLI to set the maximum time for guest accounts 821
Managing Files on the Controller 821
Transferring ArubaOS Image Files 822
In the WebUI 822
In the CLI 822
Backing Up and Restoring the Flash File System 823
Backup the Flash File System in the WebUI 823
Backup the Flash File System in the CLI 823
Restore the Flash File System in the WebUI 823
Restore the Flash File System in the CLI 823
Copying Log Files 823
In the WebUI 823
In the CLI 823
Copying Other Files 824
In the WebUI 824
In the CLI 824
Setting the System Clock 824
Manually Setting the Clock 824
In the WebUI 824
In the CLI 824

Clock Synchronization 825
In the WebUI 825
In the CLI 825
Configuring NTP Authentication 825
In the WebUI 825
In the CLI 826
Timestamps in CLI Output 826
ClearPass Profiling with IF-MAP 826
In the WebUI 826
In the CLI 826
Whitelist Synchronization 827
In the WebUI 827
In the CLI 827
Downloadable Regulatory Table 828
Important Points to Remember 828
Copying the Regulatory-Cert 828
In the WebUI 829
In the CLI 829
Activating the Regulatory-Cert 829
In the WebUI 829
In the CLI 829
Related Show Commands 830
802.11u Hotspots 831
Hotspot 2.0 Pre-Deployment Information 831
Hotspot Profile Configuration Tasks 831
Hotspot 2.0 Overview 831
Generic Advertisement Service (GAS) Queries 831
ANQP Information Elements 832

Hotspot Profile Types 832
Configuring Hotspot 2.0 Profiles 834
In the WebUI 834
In the CLI 838
Configuring Hotspot Advertisement Profiles 839
Configuring an Advertisement Profile 839
In the WebUI 839
In the CLI 840
Associating the Advertisement Profile to a Hotspot 2.0 Profile 840
In the WebUI 840
In the CLI 840
Configuring ANQP Venue Name Profiles 840
In the WebUI 841
Venue Types 841
In the CLI 842
Configuring ANQP Network Authentication Profiles 842
In the WebUI 842
In the CLI 843
Configuring ANQP Domain Name Profiles 843
In the WebUI 843
In the CLI 844
Configuring ANQP IP Address Availability Profiles 844
In the WebUI 844
In the CLI 845
Configuring ANQP NAI Realm Profiles 845
In the WebUI 845
In the CLI 848
Configuring ANQP Roaming Consortium Profiles 848
In the WebUI 848

In the CLI 849
Configuring ANQP 3GPP Cellular Network Profiles 849
In the WebUI 849
In the CLI 850
Configuring H2QP Connection Capability Profiles 850
In the WebUI 850
In the CLI 851
Configuring H2QP Operator Friendly Name Profiles 852
In the WebUI 852
In the CLI 852
Configuring H2QP Operating Class Indication Profiles 853
In the WebUI 853
In the CLI 853
Configuring H2QP WAN Metrics Profiles 853
In the WebUI 853
In the CLI 854
Adding Local Controllers 857
Configuring Local Controllers 857
Using the Initial Setup 857
Using the Web UI 857
Using the CLI 858
Configuring Layer-2/Layer-3 Settings 858
Configuring Trusted Ports 858
Configuring Local Controller Settings 858
Configuring APs 859
Using the WebUI to configure the LMS IP 859
Using the CLI to configure the LMS IP 859
Moving to a Multi-Controller Environment 859

Configuring a Preshared Key 860
Using the WebUI to configure a Local Controller PSK 860
Using the WebUI to configure a Master Controller PSK 861
Using the CLI to configure a PSK 861
Master Controller 861
Local Controller 861
Configuring a Controller Certificate 861
Using the CLI to configure a Local Controller Certificate 861
Using the CLI to configure the Master Controller Certificate 862
Advanced Security 863
Securing Client Traffic 863
Securing Wireless Clients 864
In the WebUI 864
In the CLI 865
Securing Wired Clients 865
In the WebUI 866
In the CLI 867
Securing Wireless Clients Through Non-Dell APs 867
In the WebUI 867
In the CLI 868
Securing Clients on an AP Wired Port 868
In the WebUI 868
In the CLI 869
Enabling or Disabling the Spanning Tree Parameter in AP Wired Port Profile 870
Using the WebUI 870
Using the CLI 870
Securing Controller-to-Controller Communication 870
Configuring Controllers for xSec 871
In the WebUI 871

In the CLI 871
Configuring the Odyssey Client on Client Machines 871
Installing the Odyssey Client 871
Voice and Video 878
Voice and Video License Requirements 878
Configuring Voice and Video 878
Voice ALG and Network Address Translation 878
Setting up Net Services 878
Using Default Net Services 878
Creating Custom Net Services 879
Configuring User Roles 879
Using the Default User Role 879
Creating or Modifying Voice User Roles 880
Using the WebUI to configure user roles 880
Using the CLI to configure a user role 881
Using the User-Derivation Roles 882
Using the WebUI to Derive the Role Based on SSID 882
Using the CLI to Derive the Role Based on SSID 882
Using the WebUI to Derive the Role Based on MAC OUI 882
Using the CLI to Derive the Role Based on MAC OUI 882
Configuring Firewall Settings for Voice and Video ALGs 883
In the WebUI 883
In the CLI 883
Additional Video Configurations 883
Configuring Video over WLAN enhancements 883
Prerequisites 884
In the WebUI 884
In the CLI 887
Working with QoS for Voice and Video 891

Understanding VoIP Call Admission Control Profile 891
In the WebUI 891
In the CLI 892
Understanding Wi-Fi Multimedia 893
Enabling WMM 893
In the WebUI 893
In the CLI 893
Configuring WMM AC Mapping 894
Using the WebUI to map between WMM AC and DSCP 894
Using the CLI to map between WMM AC and DSCP 895
Configuring DSCP Priorities 895
Configuring Dynamic WMM Queue Management 896
Enhanced Distributed Channel Access 896
Using the WebUI to configure EDCA parameters 897
Using the CLI to configure EDCA parameters 898
Enabling WMM Queue Content Enforcement 899
In the WebUI 899
In the CLI 899
Unified Communication and Collaboration 899
Microsoft® Lync Visibility and Granular QoS Prioritization 899
Lync ALG Compatibility Matrix 900
Configuration Prerequisites 901
Configuring Lync ALG 901
Configuring Lync Listening Port 901
Configuring Lync ALG Status 902
Dynamically Open Firewall for UCC Clients using STUN 902
Configuring Per User Role Lync Call Prioritization 903
Disable Media Classification 905
Viewing Lync ALG Statistics using the CLI 905
Viewing the list of Lync Clients 906

Viewing Call Detail Record for Lync Calls 906
Viewing Call Quality for Lync Calls 906
Viewing Lync Call Trace Buffer 906
Viewing Lync ALG Statistics Using the WebUI 906
Viewing Voice Status 906
Viewing Call Performance Report 906
Viewing Call Density Report 906
Viewing Call Detail Report 907
Viewing Voice Client Call Statistics 907
Viewing Voice Client HandOff Information 907
Viewing Voice Client Troubleshooting Information 907
Troubleshooting Lync ALG Issues 907
Enabling Lync ALG Debug Logs 907
Viewing Lync ALG Debug Logs 907
UCC Dashboard in the WebUI 908
UCC Dashboard Aggregated Display 908
Chart View 908
Details View 909
UCC Dashboard Per Client Display 910
Viewing UCC Information 912
Viewing UCC Call Detailed Record 912
Viewing UCC Client Information 912
Viewing UCC Configuration 912
Viewing UCC Statistics 912
Viewing UCC Trace Buffer 912
UCC Troubleshooting 912
UCC-AirWave Integration 912
UCC Call Quality Metrics 913
Changes to Call Admission Control 914
UCC Limitations 915

Understanding Extended Voice and Video Features 915
Understanding QoS for Microsoft Lync and Apple Facetime 915
Microsoft Lync 915
Microsoft Lync Support for Mobile Devices 915
Apple Facetime 916
In the WebUI 916
Enabling WPA Fast Handover 917
In the WebUI 917
In the CLI 918
Enabling Mobile IP Home Agent Assignment 918
Scanning for VoIP-Aware ARM 918
In the WebUI 918
In the CLI 918
Disabling Voice-Aware 802.1x 918
In the WebUI 919
In the CLI 919
Configuring SIP Authentication Tracking 919
In the WebUI 919
In the CLI 919
Enabling Real Time Call Quality Analysis 919
Important Points to Remember 920
In the Web UI 920
Viewing Real Time Call Quality Reports 920
In the CLI 920
Enabling SIP Session Timer 921
In the WebUI 921
In the CLI 922
Enabling Wi-Fi Edge Detection and Handover for Voice Clients 922
In the WebUI 922

In the CLI 923
Working with Dial Plan for SIP Calls 923
Understanding Dial Plan Format 923
Configuring Dial Plans 924
In the WebUI 924
In the CLI 926
Enabling Enhanced 911 Support 926
Working with Voice over Remote Access Point 927
Understanding Battery Boost 928
In the WebUI 928
In the CLI 928
Enabling LLDP 929
In the WebUI 929
In the CLI 932
Advanced Voice Troubleshooting 933
Viewing Troubleshooting Details on Voice Client Status 933
In the WebUI 934
In the CLI 934
Viewing Troubleshooting Details on Voice Call CDRs 935
In the WebUI 935
In the CLI 936
Enabling Voice Logs 936
In the WebUI 936
Enabling Logging for a Specific Client 937
In the CLI 937
Viewing Voice Traces 937
In the WebUI 937
In the CLI 937
Viewing Voice Configurations 938

In the CLI 938

AirGroup 940

Zero Configuration Networking 940

AirGroup Solution 940

AirGroup Services 941

AirGroup Solution Components 942

AirGroup and ClearPass Policy Manager 942

AirGroup Deployment Models 944

Integrated Deployment Model 944

AirGroup with ClearPass Policy Manager 945

Features Supported in AirGroup 945

Multi-Controller AirGroup Cluster 945

Multi-Controller AirGroup Cluster—Terminologies 945

AirGroup Domain 945

AirGroup Cluster 945

Active-Domain 946

Sample AirGroup Cluster Topology 946

Domain Definition 947

Active-Domain Definition 947

AirGroup Controller Communication 947

AirGroup Server Discovery 947

Scalability 948

Master-Local Controller Synchronization 948

Pre-configured AirGroup Services 948

AirGroup IPv6 Support 949

Limitations 949

DLNA UPnP Support 949

AirGroup mDNS Static Records 949

Group Based Device Sharing 949

Dashboard Monitoring Enhancements 949
ClearPass Policy Manager and ClearPass Guest Features 950
Best Practices and Limitations 950
Apple iTunes Wi-Fi Synchronization and File Sharing 950
Firewall Configuration 950
Disable Inter-User Firewall Settings 950
ValidUser ACL Configuration 950
Allow GRE and UDP 5353 950
Recommended Ports 951
Ports for AirPlay Service 951
Ports for AirPrint Service 951
AirGroup Services for Large Deployments 951
AirGroup Scalability Limits 952
Memory Utilization 952
CPU Utilization 952
General AirGroup Limitations 953
Integrated Deployment Model 954
Master-Local Controller Synchronization 954
Configuring an AirGroup Integrated Deployment Model 955
Enabling or Disabling AirGroup Global Setting 955
Using the WebUI 955
Using the CLI 956
Enabling or Disabling mDNS and DLNA 956
Using the CLI 956
Viewing AirGroup Global Setting on Controller 956
Using the WebUI 956
Using the CLI 956
Defining an AirGroup Service 957
Using the WebUI 957

Using the CLI 958
Enabling the allowall Service 959
Using the WebUI 959
Using the CLI 959
Enabling or Disabling an AirGroup Service 959
Using the WebUI 959
Using the CLI 960
Viewing AirGroup Service Status 960
Using the WebUI 960
Using the CLI 960
Viewing Blocked Services 960
Using the CLI 960
Viewing AirGroup Service Details 960
Using the WebUI 960
Using the CLI 960
Configuring an AirGroup Domain 960
Using the WebUI 961
Using the CLI 961
Viewing an AirGroup Domain 961
Using the WebUI 961
Using the CLI 961
Configuring an AirGroup active-domain 961
Using the WebUI 961
Using the CLI 962
Viewing an AirGroup active-domains 962
Using the WebUI 962
Using the CLI 962
Viewing AirGroup VLAN Table 962
Using the WebUI: 962
Using the CLI 962
Viewing AirGroup Multi-Controller Table 962

Using the CLI 962
Controller Dashboard Monitoring 962
Configuring the AirGroup-CPPM Interface 965
Configuring the CPPM Query Interval 965
Using the WebUI 965
Using the CLI 966
Viewing the CPPM Query Interval 966
Using the WebUI 966
Using the CLI 966
Defining a CPPM and RFC3576 Server 966
Configuring a CPPM Server 967
Using the WebUI 968
Using the CLI 968
Configuring the CPPM Server Group 968
Using the WebUI 968
Using the CLI 968
Configuring an RFC 3576 Server 968
Using the WebUI 968
Using the CLI 969
Assigning CPPM and RFC 3576 Servers to AirGroup 969
Using the WebUI 969
Using the CLI 969
Sample Configuration 970
Viewing the CPPM Server Configuration 970
Using the WebUI 970
Using the CLI 970
Verifying CPPM Device Registration 970
Configuring CPPM to Enforce Registration 971
Using the WebUI 971
Using the CLI 971

Group Based Device Sharing 971
Example 972
AirGroup mDNS Static Records 973
Important Points to Remember 973
Creating mDNS Static Records on a Controller 973
Group mDNS Static Records 973
Creating a PTR Record 973
Creating an SRV Record 974
Creating an A Record 974
Creating an AAAA Record 974
Creating a TEXT Record 974
Individual Static mDNS Records 974
Creating an Individual SRV Record 974
Creating an Individual TEXT Record 974
Creating an Individual A Record 974
Creating an Individual AAAA Record 974
Troubleshooting and Log Messages 975
Controller Troubleshooting Steps 975
ClearPass Guest Troubleshooting Steps 975
ClearPass Policy Manager Troubleshooting Steps 975
Log Messages 975
Show Commands 976
Viewing AirGroup mDNS and DLNA Cache 976
Viewing AirGroup mDNS and DLNA Statistics 976
Viewing AirGroup VLANs 976
Viewing AirGroup Servers 976
Viewing AirGroup Users 976
Viewing Service Queries Blocked by AirGroup 976
Viewing Blocked Services 976
AirGroup Global Tokens 976

Instant AP VPN Support 978
Overview 978
Improved DHCP Pool Management 978
Termination of Instant AP VPN Tunnels 978
Termination of IAP GRE Tunnels 978
L2/L3 Network Mode Support 979
Instant AP VPN Scalability Limits 979
Instant AP VPN OSPF Scaling 979
Branch-ID Allocation 981
Centralized BID Allocation 981
VPN Configuration 982
Whitelist DB Configuration 982
Controller Whitelist DB 982
External Whitelist DB 982
VPN Local Pool Configuration 982
Role Assignment for the Authenticated IAPs 982
VPN Profile Configuration 983
Viewing Branch Status 983
Example 983
W-600 Series Controllers 985
Connecting with a USB Cellular Modems 985
How it Works 985
Switching Modes 985
Finding USB Modem Commands 986
Uplink Manager 986
Cellular Profile 986
Dialer Group 987
Configuring a Supported USB Modem 987

Configuring a New USB Modem 988
Configuring the Profile and Modem Driver 989
Configuring the TTY Port 989
Testing the TTY Port 990
Selecting the Dialer Profile 991
Linux Support 991
External Services Interface 992
Sample ESI Topology 992
Understanding the ESI Syslog Parser 994
ESI Parser Domains 994
Peer Controllers 995
Syslog Parser Rules 996
Condition Pattern Matching 996
User Pattern Matching 996
Configuring ESI 996
Configuring Health-Check Method, Groups, and Servers 997
In the WebUI 997
In the CLI 997
Defining the ESI Server 998
In the WebUI 998
In the CLI 998
Defining the ESI Server Group 998
In the WebUI 998
In the CLI 999
Redirection Policies and User Role 999
In the WebUI 999
In the CLI 1000
ESI Syslog Parser Domains and Rules 1000

Managing Syslog Parser Domains in the WebUI 1000
Adding a new syslog parser domain 1000
Deleting an existing syslog parser domain 1001
Editing an existing syslog parser domain 1001
Managing Syslog Parser Domains in the CLI 1001
Adding a new syslog parser domain 1001
Showing ESI syslog parser domain information 1001
Deleting an existing syslog parser domain 1001
Editing an existing syslog parser domain 1001
Managing Syslog Parser Rules 1002
In the WebUI 1002
Adding a new parser rule 1002
Deleting a syslog parser rule 1002
Editing an existing syslog parser rule 1003
Testing a Parser Rule 1003
In the CLI 1003
Adding a new parser rule 1003
Showing ESI syslog parser rule information: 1004
Deleting a syslog parser rule: 1004
Editing an existing syslog parser rule 1004
Testing a parser rule 1004
Monitoring Syslog Parser Statistics 1004
In the WebUI 1004
In the CLI 1004
Sample Route-mode ESI Topology 1004
ESI server configuration on controller 1005
IP routing configuration on Fortinet gateway 1005
Configuring the Example Routed ESI Topology 1005
Health-Check Method, Groups, and Servers 1006

Defining the Ping Health-Check Method 1006
In the WebUI 1006
In the CLI 1006
Defining the ESI Server 1006
In the WebUI 1006
In the CLI 1007
Defining the ESI Server Group 1007
In the WebUI 1007
In the CLI 1007
Redirection Policies and User Role 1008
In the WebUI 1008
In the CLI 1008
Syslog Parser Domain and Rules 1009
Add a New Syslog Parser Domain in the WebUI 1009
Adding a New Parser Rule in the WebUI 1009
In the CLI 1010
Sample NAT-mode ESI Topology 1010
ESI server configuration on the controller 1012
Configuring the Example NAT-mode ESI Topology 1012
Configuring the NAT-mode ESI Example in the WebUI 1012
In the WebUI 1012
Configuring the ESI Group in the WebUI 1013
Configure the ESI Servers in the WebUI 1013
Configuring the Redirection Filter in the WebUI 1013
Configuring the Example NAT-mode Topology in the CLI 1014
Configuring a Health-Check Ping 1014
Configuring ESI Servers 1014
Configure an ESI Group, Add the Health-Check Ping and ESI Servers 1015
Using the ESI Group in a Session Access Control List 1015

CLI Configuration Example 1 1015
CLI Configuration Example 2 1015
Understanding Basic Regular Expression (BRE) Syntax 1016
Character-Matching Operators 1016
Regular Expression Repetition Operators 1016
Regular Expression Anchors 1017
References 1017
External User Management 1019
Overview 1019
Before you Begin 1019
Working with the ArubaOS XML API Works 1019
Creating an XML Request 1019
Adding a User 1020
Deleting a User 1020
Authenticating a User 1020
Blacklisting a User 1021
Querying for User Status 1021
XML Response 1021
Default Response Format 1021
Response Codes 1022
Query Command Response Format 1023
Using the XML API Server 1024
Configuring the XML API Server 1024
Associating the XML API Server to a AAA profile 1025
Set up Captive Portal profile 1026
Associating the Captive Portal Profile to an Initial Role 1027
Creating an XML API Request 1027
Monitoring External Captive Portal Usage Statistics 1028

Sample Code 1029
Using XML API in C Language 1029
Understanding Request and Response 1032
Understanding XML API Request Parameters 1032
Understanding XMl API Response 1033
Adding a Client 1033
Response from the controller 1034
View the updated details of the client on the controller 1034
Deleting a Client 1034
Response from the controller 1034
Authenticating a Client 1035
Status of the client before authentication 1035
Sending the authentication command 1035
Response from the controller 1035
Status of the client after authentication 1036
Querying for Client Details 1036
Response from the controller 1036
Blacklisting a Client 1037
Response from the controller 1037
Behavior and Defaults 1039
Understanding Mode Support 1039
Understanding Basic System Defaults 1040
Network Services 1040
Policies 1042
Validuser and Logon-control ACLs 1046
Roles 1047
Understanding Default Management User Roles 1049
Understanding Default Open Ports 1052

DHCP with Vendor-Specific Options 1055
Configuring a Windows-Based DHCP Server 1055
Configuring Option 60 1055
To configure option 60 on the Windows DHCP server 1055
Configuring Option 43 1056
To configure option 43 on the Windows DHCP server: 1056
Enabling DHCP Relay Agent Information Option (Option 82) 1058
Configuring Option 82 1058
In the WebUI 1058
In the CLI 1058
Enabling Linux DHCP Servers 1059
802.1X Configuration for IAS and Windows Clients 1060
Configuring Microsoft IAS 1060
RADIUS Client Configuration 1060
Remote Access Policies 1060
Active Directory Database 1061
Configuring Policies 1061
Configuring RADIUS Attributes 1062
Configuring Management Authentication using IAS 1062
Creating a Remote Policy 1063
Defining Properties for Remote Policy 1063
Creating a User Entry in Windows Active Directory 1063
Configure the Controller to use IAS Management Authentication 1064
Verify Communication between the Controller and the RADIUS Server 1064
Window XP Wireless Client Sample Configuration 1064
Acronyms and Terms 1067
Acronyms 1067
Terms 1073

This User Guide describes the features supported by Dell Networking W-Series ArubaOS 6.4.x and provides instructions and examples for configuring Dell mobility controllers and access points (APs). This guide is intended for system administrators responsible for configuring and maintaining wireless networks and assumes administrator knowledge in Layer 2 and Layer 3 networking technologies.

This chapter covers the following topics:

- What's New In ArubaOS 6.4.x on page 85
- Fundamentals on page 91
- Related Documents on page 92
- Conventions on page 92
- Related Documents on page 92

What's New In ArubaOS 6.4.x

The following features are introduced or enhanced in ArubaOS 6.4.2.0:

Table 1: *New Features/Enhancements in ArubaOS 6.4.2.0*

Feature	Description
Enhanced LACP support on W-AP220 Series and W-AP270 Series access points	This enhanced LACP feature allows W-AP220 Series or W-AP270 Series access points to form a 802.11g radio tunnel to a backup controller in the event of a controller failover, even if the backup controller is in a different L3 network.
RTLS Station Message Frequency	Currently, when configuring the RTLS server under ap system-profile, the valid range of values for **station-message-frequency** was 5-3600 seconds. There are deployments that may require this to be configurable to as frequently as 1 per second. Starting with ArubaOS 6.4.2.0, you can set the **station-message-frequency** parameter in the 1-3600 seconds range.
Service Tag	A service tag is a unique seven digit alphanumeric string that is used to electronically identify a Dell device. It is similar to a serial number identifier. Starting with ArubaOS 6.4.2.0, you can view the service tag of some newer Dell APs from the controller WebUI or CLI. It is displayed along with the serial number in a device information listing.
VHT Support on W-AP200 Series, W-AP210 Series, W-AP220 Series, and W-AP270 Series Access Points	Starting with ArubaOS 6.4.2.0, VHT is supported on W-AP220 Series access points on both 20 MHz and 40 MHz channels.
Web Content Classification	The WebCC feature helps classify web traffic in the controller. The classification is done in the data path while the traffic flows through the controller and updates dynamically. WebCC uses a cloud-based service to dynamically determine the types of websites being visited, and their safety.

Table 2: *New Hardware Platforms in ArubaOS 6.4.2.0*



Check with your local Dell sales representative on new controllers and access points availability in your country.

Hardware	Description
W-AP210 Series	The Dell W-AP210 Series (W-AP214 and W-AP215) wireless access points support the IEEE 802.11ac standard for high-performance WLAN. These access points use MIMO (Multiple-Input, Multiple-Output) technology and other high-throughput mode techniques to deliver high-performance, 802.11ac 2.4 GHz and 802.11ac 5 GHz functionality while simultaneously supporting existing 802.11a/b/g wireless services. The W-AP210 Series access points work only in conjunction with a Dell controller. The Dell W-AP210 Series access point provides the following capabilities: • Wireless transceiver • Protocol-independent networking functionality • IEEE 802.11a/b/g/n/ac operation as a wireless access point • IEEE 802.11a/b/g/n/ac operation as a wireless air monitor • Compatibility with IEEE 802.3at PoE+ and 802.3af PoE • Central management configuration and upgrades through a controller For more information, see the *Dell Networking W-AP210 Series Wireless Access Point Installation Guide*.

The following features are introduced or enhanced in ArubaOS 6.4.1.0:

Table 3: *New Features/Enhancements in ArubaOS 6.4.1.0*

Feature	Description
AirGroup	The following AirGroup service changes are effective in this release: • The **Chromecast** service is renamed to **DIAL**. • The **googlecast** service is introduced.
AP Fast Failover support for Bridge-mode Virtual AP	High Availability (HA) support for bridge mode in Campus AP is introduced in this release. In previous versions of ArubaOS the fast failover feature for Campus AP was supported using tunnel or decrypt mode. Now support has been extended to bridge mode as well.
Authentication Profile based User Idle Timeout	The **user-idle-timeout** parameter under AAA profile accepts a value of 0. Entering a value of 0, L3 user state is removed immediately upon disassociation. In other words, the controller deletes the user immediately after disassociation or disconnection from the wireless network. If RADIUS accounting is configured, the controller sends an accounting STOP message to the RADIUS server. **NOTE:** User idle timeout of 0 should not be configured for wired, split-tunnel, VIA, and VPN users. It is applicable only for wireless users in tunnel and decrypt-tunnel forwarding modes.

Table 3: *New Features/Enhancements in ArubaOS 6.4.1.0*

Feature	Description
DHCP Lease Limit	This section outlines the maximum number of DHCP leases supported for the new W-7000 Series controller platform.
Downloadable Regulatory Table	The downloadable regulatory table features allows new regulatory approvals to be distributed without waiting for a new software patch and upgrade. A separate file, called the Regulatory-Cert, containing AP regulatory information will be released periodically and made available for download on the customer support site. The Regulatory-Cert file can then be uploaded to a controller and pushed to deployed APs.
Global Firewall Parameters	The following new parameters are introduced: • **Monitor/police ARP attack (non Gratuitous ARP) rate (per 30 sec)** • **Monitor/police Gratuitous ARP attack rate (per 30 sec)**

Table 4: *New Hardware Platforms in ArubaOS 6.4.1.0*



Check with your local Dell sales representative on new controllers and access points availability in your country.

Hardware	Description
W-7000 Series Controllers	The Dell W-7000 Series controllers is an integrated controller platform. The platform acts as a software services platform targeting small to medium branch offices and enterprise networks. The W-7000 Series controllers include three models that provide varying levels of scalability. • W-7005 • W-7010 • W-7030 For more information, see the installation guide for each controller model.
W-AP103H	The Dell W-AP103H wireless access point supports the IEEE 802.11n standard for high-performance WLAN. It is a dual radio, 2x2:2 802.11n access point. This access point uses MIMO (Multiple-Input, Multiple-Output) technology and other high-throughput mode techniques to deliver high-performance 802.11n 2.4 GHz or 5 GHz functionality while simultaneously supporting existing 802.11a/b/g wireless services. The W-AP103H access point is equipped with a total of three active Ethernet ports (ENET 0-2). It is a wall-box type access point. The W-AP103H access point works only with a Dell controller. The Dell W-AP103H access point provides the following capabilities: • Wireless transceiver • Protocol-independent networking functionality • IEEE 802.11a/b/g/n operation as a wireless access point • IEEE 802.11a/b/g/n operation as a wireless air monitor • Compatibility with IEEE 802.3af PoE • Central management configuration and upgrades through a controller For more information, see the *Dell Networking W-AP103H Wireless Access Point Installation Guide*.
W-AP200 Series	The Dell W-AP200 Series (W-AP204 and W-AP205) wireless access points support the IEEE 802.11ac and 802.11n standards for high-performance WLAN. It is a dual radio, 2x2:2 802.11ac access point. These access points use MIMO (Multiple-Input, Multiple-Output) technology and other high-throughput mode techniques to deliver high-performance, 802.11n 2.4 GHz and 802.11ac 5 GHz functionality while simultaneously supporting legacy 802.11a/b/g wireless services. The Dell W-AP200 Series access point provides the following capabilities: • Wireless transceiver • Protocol-independent networking functionality • IEEE 802.11a/b/g/n/ac operation as a wireless access point • IEEE 802.11a/b/g/n/ac operation as a wireless air monitor • Compatibility with IEEE 802.3af PoE • Central management configuration and upgrades through a controller For more information, see the *Dell Networking W-AP200 Series Wireless Access Point Installation Guide*.

What's New In ArubaOS 6.4.0.0

The following features are introduced in ArubaOS 6.4.0.0:

Table 5: *New Features in ArubaOS 6.4.0.0*

Feature	Description
W-AP270 Series Access Points	The DellW-AP270 Series (W-AP274 and W-AP275) wireless access points are environmentally hardened, outdoor rated, dual-radio IEEE 802.11ac wireless access points. These access points use MIMO (Multiple-Input, Multiple-Output) technology and other high-throughput mode techniques to deliver high-performance, 802.11ac 2.4 GHz and 5 GHz functionality while simultaneously supporting existing 802.11a/b/g/n wireless services.
W-AP103 Access Point	The W-AP103 wireless access point supports the IEEE 802.11n standard for high-performance WLAN. This access point uses MIMO (Multiple-Input, Multiple-Output) technology and other high-throughput mode techniques to deliver high performance, 802.11n 2.4 GHz or 5 GHz functionality while simultaneously supporting existing 802.11a/b/g wireless services.
Ability to Disable Factory-Default IKE/IPsec Profiles	This feature enables you to disable default IKE policies, default IPsec dynamic maps, and site-to-site IPsec maps.
AirGroup	The AirGroup feature has been enhanced with the following new features in ArubaOS 6.4: • DLNA UPnP support • Group Based Device Sharing • AirGroup mDNS Static Records • Dashboard Monitoring Enhancements
Application Single Sign-On Using Layer 2 Authentication Information	This feature allows single sign-on for web-based applications using layer 2 authentication information. With single sign-on, a user does not need to provide authentication credentials before logging into each application.
AppRF 2.0	This feature improves application visibility and control by allowing you to configure and view access control list (ACL) and bandwidth application and application category-specific data. AppRF 2.0 supports a Deep Packet Inspection (DPI) engine for application detection for over a thousand applications.
AppRF Application Dashboard Visibility	This feature is supported only in the W-7200 Series controller. This feature allows you to configure both application and application category policies within a given user role. The AppRF page displays the PEF summary of all the sessions in the controller aggregated by users, devices, destinations, applications, WLANs, and roles.The elements are now represented in box charts instead of pie charts.
Authentication Server Load Balancing	Load balancing of authentication servers ensures that the authentication load is split across multiple authentication servers, thus avoiding any one particular authentication server from being overloaded.
Centralized BID Allocation	In a Master-Local set-up, the Master controller runs the BID allocation algorithm to allocate BID to the branches terminating on it and to the Local controller.
GRE Tunnels	Static IPv6 L2/L3 GRE tunnels can now be established between Dell devices and other devices that support IPv6 GRE tunnel.

Table 5: *New Features in ArubaOS 6.4.0.0*

Feature	Description
Multicast Listener Discovery	The Source Specific Multicast (SSM) option supports delivery of multicast packets that originate only from a specific source address requested by the receiver.
Hotspot 2.0	Hotspot 2.0 is a Wi-Fi Alliance Passpoint specification based upon the 802.11u protocol that provides wireless clients with a streamlined mechanism to discover and authenticate to suitable networks, and allows mobile users the ability to roam between partner networks without additional authentication.
IGMPv3 Support	ArubaOS 6.4 supports IGMPv3 functionality that makes Dell controllers aware of the Source Specific Multicast (SSM) and is used to optimize bandwidth of the network
Controller LLDP Support	ArubaOS 6.4 provides support for Link Layer Discovery Protocol (LLDP) on the controllers to advertise identity information and capabilities to other nodes on the network, and store the information discovered about the neighbors.
ClearPass Policy Manager Integration	ArubaOS now supports downloadable roles. By using this feature, when CPPM successfully authenticates a user, the user is assigned a role by CPPM and if the role is not defined on the controller, the role attributes can also be automatically downloaded.
High Availability	The high availability feature has been enhanced with the following new features in ArubaOS 6.4: • High Availability Configuration Using the WebUI • Extended Standby Controller Capacity • High Availability State Synchronization • High Availability Inter-controller Heartbeats
ArubaOS and ClearPass Guest Login URL Hash option	This feature enhances the security for the ClearPass Guest login URL. A new parameter called "url_hash_key"in the Captive Portal profile provides ClearPass the ability to trust and ensure that the client MAC address in the redirect URL has not been tampered with by anyone.
Palo Alto Networks Firewall Integration	This feature takes advantage of the User-Identification (User-ID) feature of the Palo Alto Networks (PAN) firewall allows network administrators to configure and enforce firewall policies based on user and user groups. User-ID identifies the user on the network based on the IP address of the device which the user is logged into. Additionally, firewall policy can be applied based on the type of device the user is using to connect to the network. Since the Dell controller maintains the network and user information of the clients on the network, it is the best source to provide the information for the User-ID feature on the PAN firewall.

Table 5: *New Features in ArubaOS 6.4.0.0*

Feature	Description
RADIUS Accounting on Multiple Servers	ArubaOS provides support for the controllers to send RADIUS accounting to multiple RADIUS servers. The controller notifies all the RADIUS servers to track the status of authenticated users. Accounting messages are sent to all the servers configured in the server group in a sequential order.
Unified Communication and Collaboration	The following new features are introduced in ArubaOS 6.4: • Per User Role Lync Call Prioritization • UCC Dashboard in the WebUI • UCC show Commands • UCC-AirWave Integration • Dynamically Open Firewall for UCC Clients using STUN • UCC Call Quality Metrics • Changes to Call Admission Control
802.11w Support	ArubaOS supports the IEEE 802.11w standard, also known as Management Frame Protection (MFP). MFP makes it difficult for an attacker to deny service by spoofing Deauth and Disassoc management frames. MFP uses 802.11i (Robust Security Network) framework that establishes encryption keys between the client and AP.

Fundamentals

Configure your controller and AP using either the Web User Interface (WebUI) or the command line interface (CLI).

WebUI

Each controller supports up to 320 simultaneous WebUI connections. The WebUI is accessible through a standard Web browser from a remote management console or workstation. The WebUI includes configuration wizards that step you through easy-to-follow configuration tasks. The wizards are:

- AP Wizard—basic AP configuration
- Controller Wizard—basic controller configuration
- LAN Wizard—creating and configuring new WLAN(s) associated with the “default” ap-group
- License Wizard—installation and activation of software licenses
- AirWave Wizard —Controllers running ArubaOS 6.3 and later can use the AirWave wizard to quickly and easily connect the controller to an AirWave server.

In addition to the wizards, the WebUI includes a Dashboard monitoring feature that provides enhanced visibility into your wireless network’s performance and usage. This allows you to easily locate and diagnose WLAN issues. For details on the WebUI Dashboard, see Dashboard Monitoring.

CLI

The CLI is a text-based interface accessible from a local console connected to the serial port on the controller or through a Telnet or Secure Shell (SSH) session.



By default, you access the CLI from the serial port or from an SSH session. You must explicitly enable Telnet on your controller in order to access the CLI via a Telnet session.

When entering commands remember that:

- commands are not case sensitive
- the space bar will complete your partial keyword
- the backspace key will erase your entry one letter at a time
- the question mark (?) will list available commands and options

Related Documents

The following guides are part of the complete documentation for the Dell user-centric network:

- *Dell Networking W-Series Controller Installation Guides*
- *Dell Networking W-Series Access Point Installation Guides*
- *Dell Networking W-Series ArubaOS Quick Start Guide*
- *Dell Networking W-Series ArubaOS User Guide*
- *Dell Networking W-Series ArubaOS Command Line Reference Guide*
- *Dell Networking W-Series ArubaOS MIB Reference Guide*
- *Dell Networking W-Series ArubaOS Release Notes*

Conventions

The following conventions are used throughout this document to emphasize important concepts:

Table 6: *Typographical Conventions*

Type Style	Description
Italics	This style is used to emphasize important terms and to mark the titles of books.
`System items`	This fixed-width font depicts the following: • Sample screen output • System prompts • Filenames, software devices, and specific commands when mentioned in the text
`Commands`	In the command examples, this bold font depicts text that you must type exactly as shown.
<*Arguments*>	In the command examples, italicized text within angle brackets represents items that you should replace with information appropriate to your specific situation. For example: `# send <text message>` In this example, you would type "send" at the system prompt exactly as shown, followed by the text of the message you wish to send. Do not type the angle brackets.
`[Optional]`	Command examples enclosed in brackets are optional. Do not type the brackets.
`{Item A | Item B}`	In the command examples, items within curled braces and separated by a vertical bar represent the available choices. Enter only one choice. Do not type the braces or bars.

The following informational icons are used throughout this guide:



Indicates helpful suggestions, pertinent information, and important things to remember.



Indicates a risk of damage to your hardware or loss of data.



Indicates a risk of personal injury or death.

Contacting Dell

Table 7: *Contact Information*

Web Site Support	
Main Website	**dell.com**
Contact Information	**dell.com/contactdell**
Support Website	**dell.com/support**
Documentation Website	**dell.com/support/manuals**

This chapter describes how to connect a Dell controller and Dell AP to your wired network. After completing the tasks described in this chapter, see Access Points (APs) on page 485 for information on configuring APs.

This chapter describes the following topics:

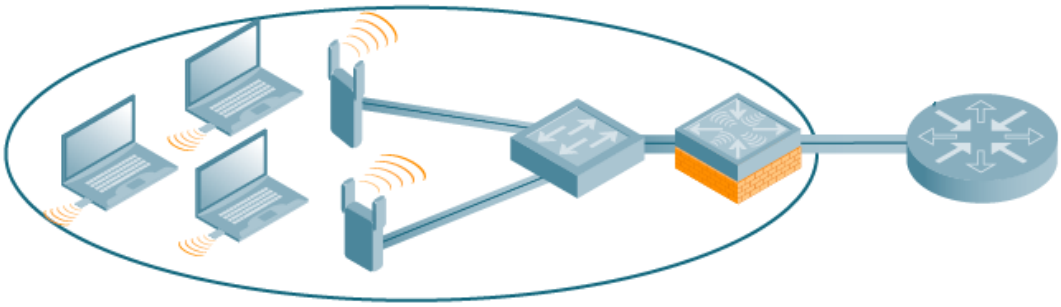
- Configuring Your User-Centric Network on page 106
- Understanding Basic Deployment and Configuration Tasks on page 94
- Configuring the Controller on page 97
- Configuring a VLAN to Connect to the Network on page 102
- Enabling Wireless Connectivity on page 105

Understanding Basic Deployment and Configuration Tasks

This section describes typical deployment scenarios and the tasks you must perform while connecting to a Dell controller and Dell AP to your wired network. For details on performing the tasks mentioned in these scenarios, refer to the other procedures within the **Basic User-Centric Networks** section of this document.

Deployment Scenario #1: Controller and APs on Same Subnet

Figure 1 *Controller and APs on Same Subnet*



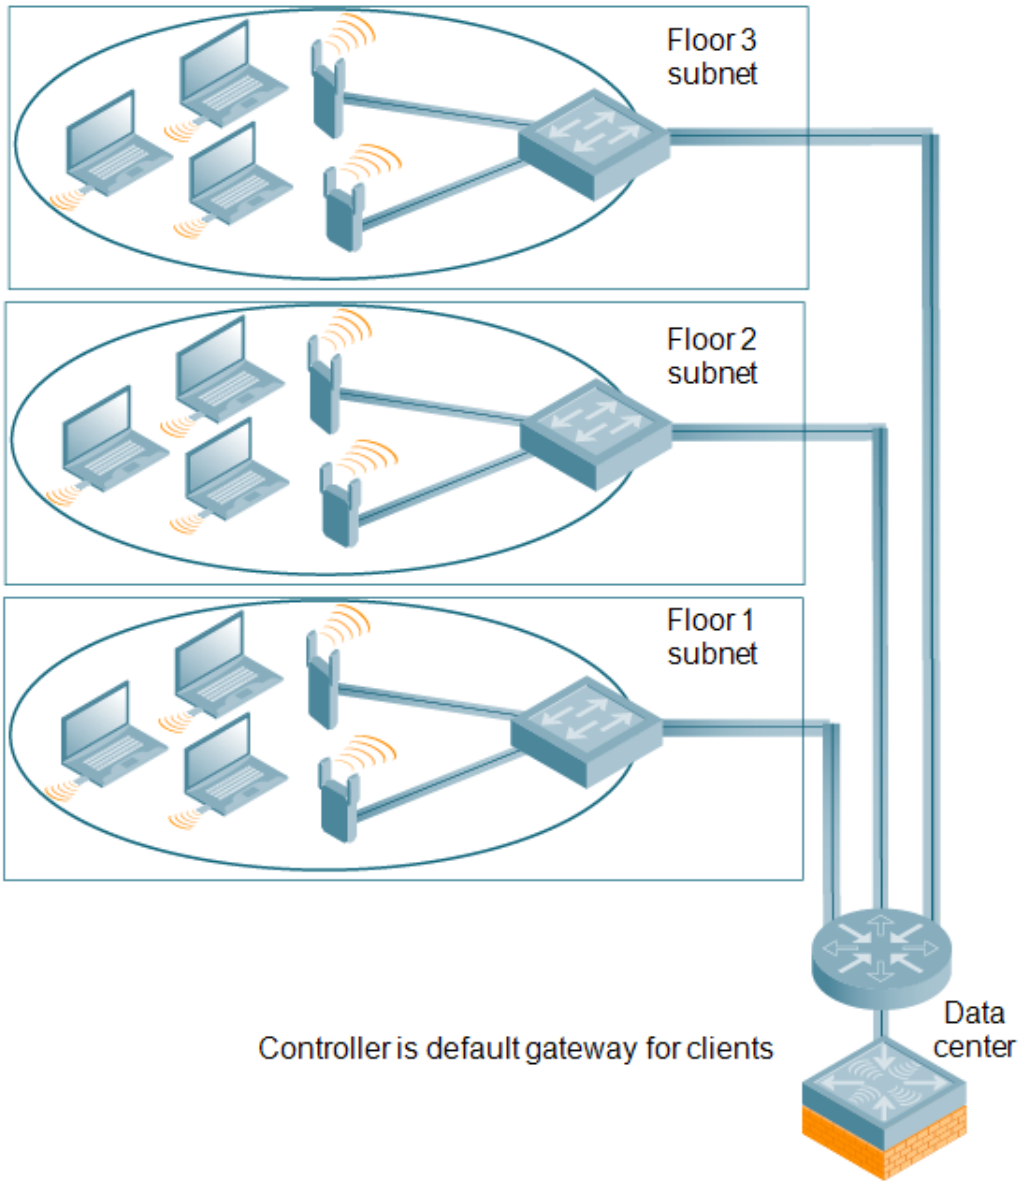
In this deployment scenario, the APs and controller are on the same subnetwork and will use IP addresses assigned to the subnetwork. The router is the default gateway for the controller and clients.There are no routers between the APs and the controller. APs can be physically connected directly to the controller. The uplink port on the controller is connected to a layer-2 switch or router.

For this scenario, you must perform the following tasks:

1. Run the initial setup wizard.
   - Set the IP address of VLAN 1.
   - Set the default gateway to the IP address of the interface of the upstream router to which you will connect the controller.
2. Connect the uplink port on the controller to the switch or router interface. By default, all ports on the controller are access ports and will carry traffic for a single VLAN.
3. Deploy APs. The APs will use the Aruba Discovery Protocol (ADP) to locate the controller.
4. Configure the SSID(s) with VLAN 1 as the assigned VLAN for all users.

Deployment Scenario #2: APs All on One Subnet Different from Controller Subnet

Figure 2 *APs All on One Subnet Different from Controller Subnets*



In this deployment scenario, the APs and the controller are on different subnetworks and the APs are on multiple subnetworks. The controller acts as a router for the wireless subnetworks (the controller is the default gateway for the wireless clients). The uplink port on the controller is connected to a layer-2 switch or router; this port is an access port in VLAN 1.

For this scenario, you must perform the following tasks:

1. Run the initial setup wizard.
   - Set the IP address for VLAN 1.
   - Set the default gateway to the IP address of the interface of the upstream router to which you will connect the controller.
2. Connect the uplink port on the controller to the switch or router interface.
3. Deploy APs. The APs will use DNS or DHCP to locate the controller.

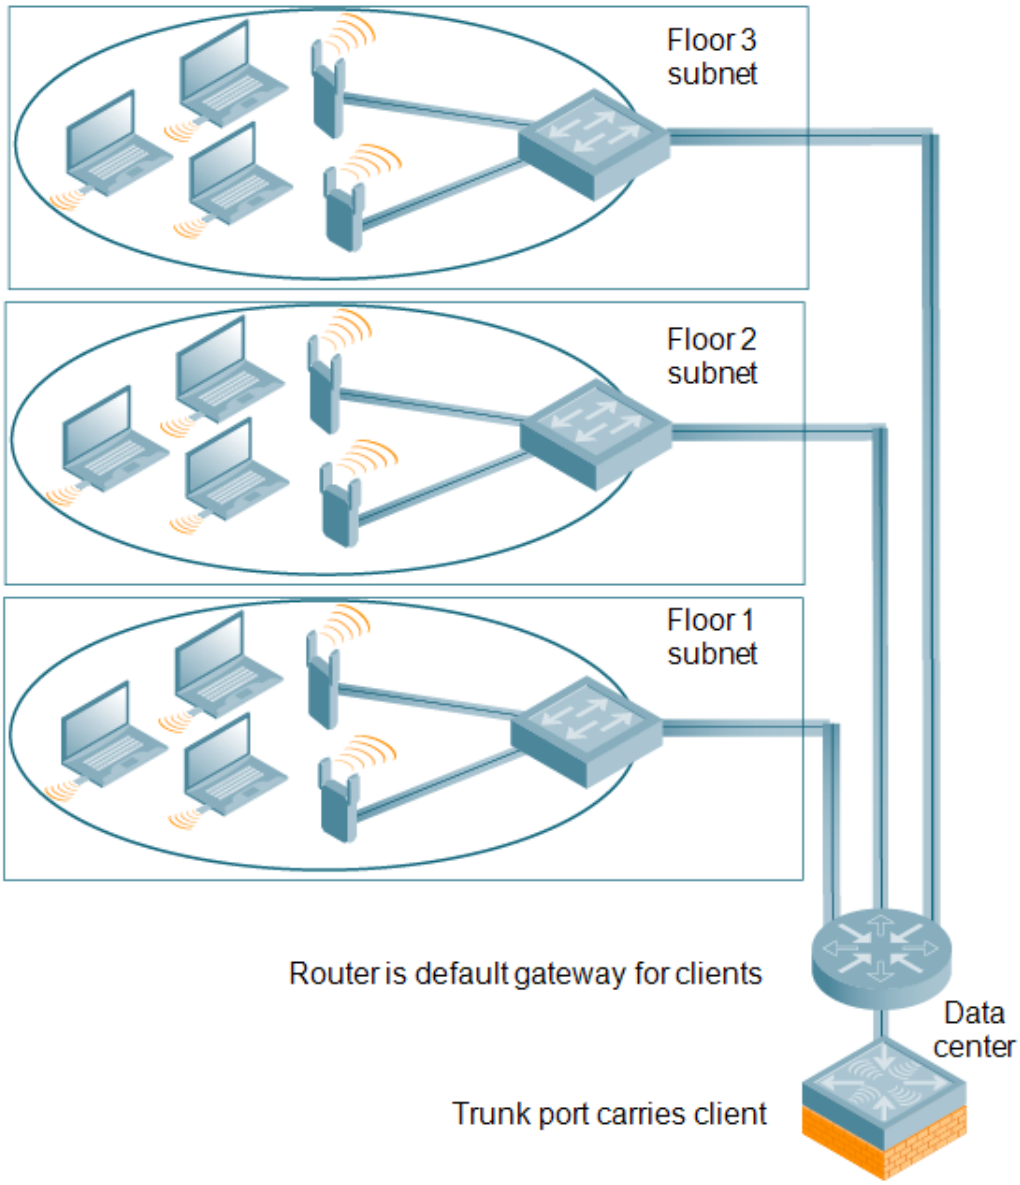
4. Configure VLANs for the wireless subnetworks on the controller.
5. Configure SSIDs with the VLANs assigned for each wireless subnetwork.



Each wireless client VLAN must be configured on the controller with an IP address. On the uplink switch or router, you must configure static routes for each client VLAN, with the controller's VLAN 1 IP address as the next hop.

Deployment Scenario #3: APs on Multiple Different Subnets from Controllers

Figure 3 *APs on Multiple Different Subnets from Controllers*



In this deployment scenario, the APs and the controller are on different subnetworks and the APs are on multiple subnetworks. There are routers between the APs and the controller. The controller is connected to a layer-2 switch or router through a trunk port that carries traffic for all wireless client VLANs. An upstream router functions as the default gateway for the wireless users.



This deployment scenario does *not* use VLAN 1 to connect to the layer-2 switch or router through the trunk port. The initial setup prompts you for the IP address and default gateway for VLAN 1; use the default values. In later steps, you configure the appropriate VLAN to connect to the switch or router as well as the default gateway.

For this scenario, you must perform the following tasks:

1. Run the initial setup.
   - Use the *default* IP address for VLAN 1. Since VLAN 1 is *not* used to connect to the layer-2 switch or router through the trunk port, you must configure the appropriate VLAN in a later step.
   - Do *not* specify a default gateway (use the default "none"). In a later step, you configure the default gateway.
2. Create a VLAN that has the same VLAN ID as the VLAN on the switch or router to which you will connect the controller. Add the uplink port on the controller to this VLAN and configure the port as a trunk port.
3. Add client VLANs to the trunk port.
4. Configure the default gateway on the controller. This gateway is the IP address of the router to which you will connect the controller.
5. Configure the loopback interface for the controller.
6. Connect the uplink port on the controller to the switch or router interface.
7. Deploy APs. The APs will use DNS or DHCP to locate the controller.
8. Now configure VLANs on the controller for the wireless client subnetworks and configure SSIDs with the VLANs assigned for each wireless subnetwork.

Configuring the Controller

The tasks in deploying a basic user-centric network fall into two main areas:

- Configuring and connecting the controller to the wired network (described in this section)
- Deploying APs (described later in this section)

To connect the controller to the wired network:

1. Run the initial setup to configure administrative information for the controller.

   Initial setup can be done using the browser-based Setup Wizard or by accessing the initial setup dialog via a serial port connection. Both methods are described in the *Dell Networking W-Series ArubaOS Quick Start Guide* and are referred to throughout this *chapter* as "initial setup."
2. (Deployment #3) Configure a VLAN to connect the controller to your network. You do *not* need to perform this step if you are using VLAN 1 to connect the controller to the wired network.
3. (Optional) Configure a loopback address for the controller. You do *not* need to perform this step if you are using the VLAN 1 IP address as the controller's IP address. Disable spanning tree on the controller if necessary.
4. Configure the system clock.
5. (Optional) Install licenses; refer to Software Licenses on page 130.
6. Connect the ports on the controller to your network.

This section describes the steps in detail.

Running Initial Setup

When you connect to the controller for the first time using either a serial console or a Web browser, the initial setup requires you to set the role (master or local) for the controller and passwords for administrator and configuration access.



Do not connect the controller to your network when running the initial setup. The factory-default controller boots up with a default IP address and both DHCP server and spanning tree functions are not enabled. Once you have completed the initial setup, you can use either the CLI or WebUI for further configuration before connecting the controller to your network.

The initial setup might require that you specify the country code for the country in which the controller will operate; this sets the regulatory domain for the radio frequencies that the APs use.



You cannot change the country code for controllers designated for certain countries, such as the U.S. Improper country code assignment can disrupt wireless transmissions. Many countries impose penalties and sanctions for operators of wireless networks with devices set to improper country codes. If none of the channels supported by the AP you are provisioning have received regulatory approval by the country whose country code you selected, the AP will revert to Air Monitor mode.

The initial setup requires that you configure an IP address for the VLAN 1 interface, which you can use to access and configure the controller remotely via an SSH or WebUI session. Configuring an IP address for the VLAN 1 interface ensures that there is an IP address and default gateway assigned to the controller upon completion of the initial setup.

Connecting to the Controller after Initial Setup

After you complete the initial setup, the controller reboots using the new configuration. (See the *Dell Networking W-Series ArubaOS Quick Start Guide* for information about using the initial setup.) You can then connect to and configure the controller in several ways using the administrator password you entered during the initial setup:

- You can continue to use the connection to the serial port on the controller to enter the command line interface (CLI). (Refer to Management Access on page 778 for information on how to access the CLI and enter configuration commands.)
- You can connect an Ethernet cable from a PC to an Ethernet port on the controller. You can then use one of the following access methods:
  - Use the VLAN 1 IP address to start an SSH session where you can enter CLI commands.
  - Enter the VLAN 1 IP address in a browser window to start the WebUI.
  - WebUi Wizards.



This chapter and the user guide in general focus on CLI and standard WebUI configuration examples. However, basic controller configuration and WLAN/LAN creation can be completed using the alternative wizards from within the WebUI. If you wish to use a configuration wizard, navigate to **Configuration > Wizards**, click on the desired wizard, and follow the imbedded help instructions within the wizard.

Dell W-7200 Series Controller

The Dell W-7200 Series controller is a new controller platform that was introduced in conjunction with ArubaOS 6.2. This controller provides new functionality and improved capabilities over previous Dell controllers. However, the W-7200 Series controller also introduces some changes that you must keep in mind when adding it to your network.

New Port Numbering Scheme

The W-7200 Series controller uses a different port numbering scheme from previous controllers. All other controller platforms use a **slot/port** numbering scheme. The W-7200 Series controller uses **slot/module/port**

instead.

It is important to consider this when migrating an older controller to the W-7200 Series. If you load a configuration from a non-W-7200 controller, that controller will not have network connectivity because any interface configuration will not be recognized. For information about migrating to a W-7200 Series controller, see the *Dell Networking W-Series ArubaOS 6.2 Release Notes*.

Individual Port Behavior

The first two ports on the W-7200 Series controller, 0/0/0 and 0/0/1 are combination ports and can be used for management, HA, and I/O. Ports 0/0/2 through 0/0/5 can only be used for I/O.

Using the LCD Screen

Some controllers are equipped with an LCD panel that displays a variety of information about the controller's status and provides a menu that allows for basic operations such as initial setup and reboot. The LCD panel displays two lines of text with a maximum of 16 characters on each line. When using the LCD panel, the active line is indicated by an arrow next to the first letter.

The LCD panel is operated using the two navigation buttons to the left of the screen.

- Menu: Allows you to navigate through the menus of the LCD panel.
- Enter: Confirms and executes the action currently displayed on the LCD panel.

The LCD has four modes:

- Boot: Displays the boot up status.
- LED Mode: Displays the mode that the STATUS LED is in.
- Status: Displays the status of different components of the controller, including Power Supplies and ArubaOS version.
- Maintenance: Allows you to execute some basic operations of the controller such as uploading an image or rebooting the system.

Table 8: *LCD Panel Mode: Boot*

Function/Menu Options	Displays
Displays boot status	"Booting ArubaOS...

Table 9: *LCD Panel Mode: LED Mode*

Function/Menu Options	Displays
Administrative	LED MODE: ADM - displays whether the port is administratively enabled or disabled.
Duplex	LED MODE: DPX - displays the duplex mode of the port.
Speed	LED MODE: SPD - displays the speed of the port.
Exit Idle Mode	EXIT IDLE MENU

Table 10: *LCD Panel Mode: Status*

Function/Menu Options	Displays
ArubaOS	Version ArubaOS X.X.X.X
PSU	Status Displays status of the power supply unit. **PSU 0: [OK | FAILED | MISSING]** **PSU 1: [OK | FAILED | MISSING]**
Fan Tray	Displays fan tray status. **FAN STATUS: [OK | ERROR | MISSING]** **FAN TEMP: [OK | HIGH | SHUTDOWN]**
Exit Status Menu	EXIT STATUS

Table 11: *LCD Panel Mode: Maintenance*

Function/Menu Options	Displays
Upgrade Image	Upgrade the software image on the selected partition from a predefined location on the attached USB flash device. **Partition [0 | 1] Upgrade Image [no | yes]**
Upload Config	Uploads the controller's current configuration to a predefined location on the attached USB flash device. **Upload Config [no | yes]**
Factory Default	Allows you to return the controller to the factory default settings. **Factory Default [no | yes]**
Media Eject	Completes the reading or writing of the attached USB device. **Media Eject [no | yes]**
System Reboot	Allows you to reboot the controller. **Reboot [no | yes]**
System Halt	Allows you to halt the controller. **Halt [no | yes]**
Exit Maintenance Menu	EXIT MAINTENANCE

Using the LCD and USB Drive

You can upgrade your image or upload your pre-saved configuration by using your USB drive and your LCD commands.

Upgrading an Image

1. Copy a new controller image onto your USB drive into a directory named **/Dellimage**.
2. Insert your USB drive into the controller's USB slot. Wait for 30 seconds for the controller to mount the USB.
3. Navigate to Upgrage Image in the LCD's Maintenance menu. Select partition and confirm the upgrade (Y/N) and then wait for controller to copy the image from USB to the system partition.
4. Execute a system reboot either from the LCD menu or from the command line to complete the upgrade.

Uploading a Pre-saved Configuration

1. Copy your pre-saved configuration and name the copied file **Dell_usb.cfg**.
2. Move your pre-saved configuration file onto your USB drive into a directory named **/Dellimage**.
3. Insert your USB drive into the controller's USB slot. Wait for 30 seconds for the controller to mount the USB.
4. Navigate to the Upload Config in the LCD's Maintenance menu. Confirm the upload (Y/N) and then wait for the upload to complete.
5. Execute a system reboot either from the LCD menu or from the command line to reload from the uploaded configuration.

For detailed upgrade and upload instruction, see the Upgrade Chapter in the Release Notes.

Disabling LCD Menu Functions

For security purpose, you can disable all LCD menu functions by disabling the entire menu functionality using the following command:

```
(host) (config) #lcd-menu
(host) (lcd-menu) #disable menu
```

To prevent inadvertent menu changes, you can disable LCD individual menu function using the following commands:

```
(host) (lcd-menu) #disable menu maintenance ?
   factory-default Disable factory default menu
   media-eject Disable media eject menu on LCD
   system-halt Disable system halt menu on LCD
   system-reboot Disable system reboot menu on LCD
   upgrade-image Disable image upgrade menu on LCD
   upload-config Disable config upload menu on LCD
```

To display the current LCD functionality from the command line, use the following command:

```
(host) (config) #show lcd-menu
lcd-menu
--------
Parameter                                   Value
---------                                   -----
menu maintenance upgrade-image partition0   enabled
menu maintenance upgrade-image partition1   enabled
menu maintenance upgrade-image              enabled
menu maintenance upload-config              enabled
menu maintenance factory-default            enabled
menu maintenance media-eject                enabled
menu maintenance reload-system              enabled
menu maintenance halt-system                enabled
menu maintenance                            enabled
menu                                        enabled
```

Configuring a VLAN to Connect to the Network

You must follow the instructions in this section only if you need to configure a trunk port between the controller and another layer-2 switch (shown in Deployment Scenario #3: APs on Multiple Different Subnets from Controllers on page 96).

This section shows how to use both the WebUI and CLI for the following configurations (subsequent steps show how to use the WebUI only):

- Create a VLAN on the controller and assign it an IP address.
- Optionally, create a VLAN pool. A VLAN pool consists of two more VLAN IDs which are grouped together to efficiently manage multi-controller networks from a single location. For example, policies and virtual application configurations map users to different VLANs which may exist at different controllers. This creates redundancy where one controller has to back up many other controllers. With the VLAN pool feature you can control your configuration globally.



VLAN pooling should *not* be used with static IP addresses.

- Assign to the VLAN the ports that you will use to connect the controller to the network. (For example, the uplink ports connected to a router are usually Gigabit ports.) In the example configurations shown in this section, a controller is connected to the network through its Gigabit Ethernet port 1/25.
- Configure the port as a trunk port.
- Configure a default gateway for the controller.

Creating, Updating, and Viewing VLANs and Associated IDs

You can create and update a single VLAN or bulk VLANS using the WebUI or the CLI. See Configuring VLANs on page 148.



In the WebUI configuration windows, clicking the **Save Configuration** button saves configuration changes so they are retained after the controller is rebooted. Clicking the **Apply** button saves changes to the running configuration but the changes are not retained when the controller is rebooted. A good practice is to use the **Apply** button to save changes to the running configuration and, after ensuring that the system operates as desired, click **Save Configuration**.

You can view VLAN IDs in the CLI.

```
(host) #configure terminal
Enter Configuration commands, one per line. End with CNTL/Z
(host) (config) #show vlan

VLAN CONFIGURATION
------------------
VLAN   Description  Ports
----   -----------  -----
1      Default      FE1/0-3 FE1/6 GE1/8
2      VLAN0002
4      VLAN0004
12     VLAN0012
210    VLAN0210
212    VLAN0212     FE1/5
213    VLAN0213     FE1/4
1170   VLAN1170     FE1/7
```

Creating, Updating, and Deleting VLAN Pools



VLAN pooling should *not* be used with static IP addresses.

You can create, update, and delete a VLAN pool using the WebUI or the CLI. See Creating a VLAN Pool on page 149.

Use the CLI to add existing VLAN IDS to a pool.

```
(host) #configure terminal
Enter Configuration commands, one per line. End with CNTL/Z
(host) (config) #vlan-name mygroup pool
(host) (config) #vlan mygroup 2,4,12
(host) (config) #
```

To confirm the VLAN pool status and mappings assignments, use the **show vlan mapping** command:

```
(host) (config) #show vlan mapping

VLAN Name      Pool Status  VLAN IDs
---------      -----------  --------
mygroup        Enabled      2,4,12
group123       Disabled
```

Assigning and Configuring the Trunk Port

The following procedures configures a Gigabit Ethernet port as trunk port.

In the WebUI

1. Navigate to the **Configuration > Network > Ports** window on the WebUI.
2. In the Port Selection section, click the port that will connect the controller to the network. In this example, click port 25.
3. For Port Mode, select **Trunk**.
4. For Native VLAN, select VLAN 5 from the scrolling list, then click the left (<--) arrow.
5. Click **Apply**.

In the CLI

```
interface gigabitethernet 1/25
   switchport mode trunk
   switchport trunk native vlan 5
```

To confirm the port assignments, use the **show vlan** command:

```
(host) (config) #show vlan

VLAN CONFIGURATION
------------------
VLAN    Name        Ports
----    ----        -----
1       Default     Fa1/0-23 Gig1/24
5       VLAN0005    Gig1/25
```

Configuring the Default Gateway

The following configurations assign a default gateway for the controller.

In the WebUI

1. Navigate to the **Configuration > Network > IP > IP Routes** window.

2. To add a new static gateway, click the **Add** button below the static IP address list.
   a. In the **IP Address** field, enter an IP address in dotted-decimal format.
   b. In the **Cost** field, enter a value for the path cost.
   c. Click **Add**.
3. You can define a dynamic gateway using DHCP, PPPOE or a cell uplink interface. In the **Dynamic** section, click the **DHCP**, **PPPoE** or **Cellular** checkboxes to select one or more dynamic gateway options. If you select more than one dynamic gateway type, you must also define a cost for the route to each gateway. The controller will first attempt to obtain a gateway IP address using the option with the lowest cost. If the controller is unable to obtain a gateway IP address, it will then attempt to obtain a gateway IP address using the option with the next-lowest path cost.
4. Click **Apply**.

In the CLI

```
ip default-gateway <ipaddr>|{import cell|dhcp|pppoe}|{ipsec <name>} <cost>
```

Configuring the Loopback IP Address for the Controller

You must configure a loopback address if you are not using a VLAN ID address to connect the controller to the network (see Deployment Scenario #3: APs on Multiple Different Subnets from Controllers on page 96).



After you configure or modify a loopback address, you must reboot the controller.

If configured, the loopback address is used as the controller's IP address. If you do not configure a loopback address for the controller, the IP address assigned to the first configured VLAN interface IP address. Generally, VLAN 1 is configured first and is used as the controller's IP address.

ArubaOS allows the loopback address to be part of the IP address space assigned to a VLAN interface. In the example topology, the VLAN 5 interface on the controller was previously configured with the IP address 10.3.22.20/24. The loopback IP address in this example is 10.3.22.220.



You configure the loopback address as a host address with a 32-bit netmask. The loopback address should be routable from all external networks.

Spanning tree protocol (STP) is enabled by default on the controller. STP ensures a single active path between any two network nodes, thus avoiding bridge loops. Disable STP on the controller if you are not employing STP in your network.

In the WebUI

1. Navigate to the **Configuration > Network > Controller > System Settings** window.
2. Enter the IP address under Loopback Interface.
3. On this window, you can also turn off spanning tree. Click **No** for Spanning Tree Enabled.
4. Click **Apply** at the bottom of the window (you might need to scroll down the window).
5. At the top of the window, click **Save Configuration**.



You must reboot the controller for the new IP address to take effect.

6. Navigate to the **Maintenance > Controller > Reboot Controller** window.

7. Click **Continue**.

In the CLI

```
interface loopback ip address 10.3.22.220
no spanning-tree
write memory
reload
```

The controller returns the following messages:

```
Do you really want to reset the system(y/n):
```

Enter **y** to reboot the controller or **n** to cancel.

```
System will now restart!
...
Restarting system.
```

To verify that the controller is accessible on the network, ping the loopback address from a workstation on the network.

Configuring the System Clock

You can manually set the clock on the controller, or configure the controller to use a Network Time Protocol (NTP) server to synchronize its system clock with a central time source. For more information about setting the controller's clock, see Setting the System Clock on page 824.

Installing Licenses

ArubaOS consists of a base operating system with optional software modules that you can activate by installing license keys. If you use the Setup Wizard during the initial setup phase, you will have the opportunity to install software licenses at that time. Refer to Software Licenses on page 130 for detailed information on Licenses.

Connecting the Controller to the Network

Connect the ports on the controller to the appropriately-configured ports on an L2 switch or router. Make sure that you have the correct cables and that the port LEDs indicate proper connections. Refer to the *Installation Guide* for the controller for port LED and cable descriptions.



In many deployment scenarios, an external firewall is situated between various Dell devices. External Firewall Configuration on page 626 describes the network ports that must be configured on the external firewall to allow proper operation of the network.

To verify that the controller is accessible on the network:

- If you are using VLAN 1 to connect the controller to the network (Deployment Scenario #2: APs All on One Subnet Different from Controller Subnet on page 95 and Deployment Scenario #3: APs on Multiple Different Subnets from Controllers on page 96), ping the VLAN 1 IP address from a workstation on the network.
- If you created and configured a new VLAN (Deployment Scenario #3: APs on Multiple Different Subnets from Controllers on page 96), ping the IP address of the new VLAN from a workstation on the network.

Enabling Wireless Connectivity

Wireless users can connect to the SSID but because you have not yet configured authentication, policies, or user roles, they will not have access to the network. Other chapters in the *Dell Networking W-Series ArubaOS*

User Guide describe how to build upon this basic deployment to configure user roles, firewall policies, authentication, authentication servers, and other wireless features.

Configuring Your User-Centric Network

Configuring your controller and AP is done through either the Web User Interface (WebUI) or the command line interface (CLI).

- WebUI is accessible through a standard Web browser from a remote management console or workstation. The WebUI includes configuration wizards that step you through easy-to-follow configuration tasks. Each wizard has embedded online help. The wizards are:
  - AP Wizard—basic AP configurations including LAN, Remote, LAN Mesh and Remote Mesh deployment scenarios
  - Controller Wizard—basic controller configuration including system settings, Control Plane security, cluster settings and licenses
  - WLAN/LAN Wizard—creating and configuring new WLANs and LANs associated with the "default" ap-group. Includes campus only and remote networking.
  - License Wizard—installation and activation of software licenses (see Software Licenses on page 130)



Clicking **Cancel** from the Wizards return you to where you launched the wizard. Any configuration changes you entered are not saved.

- The command line interface (CLI) allows you to configure and manage controllers. The CLI is accessible from a local console connected to the serial port on the controller or through a Telnet or Secure Shell (SSH) session from a remote management console or workstation.



By default, you can only access the CLI from the serial port or from an SSH session. To use the CLI in a Telnet session, you must explicitly enable Telnet on the controller.

ArubaOS supports secure IPsec communications between a controller and campus or remote APs using public-key self-signed certificates created by each master controller. The controller certifies its APs by issuing them certificates. If the master controller has any associated local controllers, the master controller sends a certificate to each local controller, which in turn sends certificates to their own associated APs. If a local controller is unable to contact the master controller to obtain its own certificate, it is not be able to certify its APs, and those APs can not communicate with their local controller until master-local communication has been reestablished. You create an initial control plane security configuration when you first configure the controller using the initial setup wizard. The ArubaOS initial setup wizard enables control plane security by default, so it is very important that the local controller be able to communicate with its master controller when it is first provisioned.

Some AP model types have factory-installed digital certificates. These AP models use their factory-installed certificates for IPsec, and do not need a certificate from the controller. Once a campus or remote AP is certified, either through a factory-installed certificate or a certificate from the controller, the AP can failover between local controllers and still stay connected to the secure network, because each AP has the same master controller as a common trust anchor.

Starting with ArubaOS 6.2, the controller maintains two separate AP whitelists; one for campus APs and one for Remote APs. These whitelists contain records of all campus APs or remote APs connected to the network. You can use a campus or AP whitelist at any time to add a new valid campus or remote AP to the secure network, or revoke network access to any suspected rogue or unauthorized APs.



The control plane security feature supports IPv4 campus and remote APs only. Do not enable control plane security on a controller that terminates IPv6 APs.

When the controller sends an AP a certificate, that AP must reboot before it can connect to its controller over a secure channel. If you are enabling control plane security for the first time on a large network, you may experience several minutes of interrupted connectivity while each AP receives its certificate and establishes its secure connection.

Topics in this chapter include:

- Control Plane Security Overview on page 107
- Configuring Control Plane Security on page 108
- Managing Whitelists on Master and Local Controllers on page 117
- Working in Environments with Multiple Master Controllers on page 120
- Replacing a Controller on a Multi-Controller Network on page 123
- Configuring Control Plane Security after Upgrading on page 127
- Troubleshooting Control Plane Security on page 128

Control Plane Security Overview

Controllers using control plane security only send certificates to APs that you have identified as valid APs on the network. If you want closer control over each AP that is certified, you can manually add individual campus and remote APs to the secure network by adding each AP's information to the whitelists when you first run the

initial setup wizard. If you are confident that all APs currently on your network are valid APs, then you can use the initial setup wizard to configure automatic certificate provisioning to send certificates from the controller to each campus or remote AP, or to all campus and remote APs within specific ranges of IP addresses.

The default automatic certificate provisioning setting requires that you manually enter each campus AP's information into the campus AP whitelist, and each remote AP's information into the remote AP whitelist. If you change the default automatic certificate provisioning values to let the controller send certificates to all APs on the network, that new setting ensures that all valid APs receive a certificate, but also increases the chance that you will certify a rogue or unwanted AP. If you configure the controller to send certificates to only those APs within a range of IP addresses, there is a smaller chance that a rogue AP receives a certificate, but any valid AP with an IP address outside the specified address ranges will not receive a certificate, and can not communicate with the controller (except to obtain a certificate). Consider both options carefully before you complete the control plane security portion of the initial setup wizard. If your controller has a publicly accessible interface, you should identify the APs on the network by IP address range. This prevents the controller from sending certificates to external or rogue campus APs that may attempt to access your controller through that publicly accessible interface.

Configuring Control Plane Security

When you initially deploy the controller, you create your initial control plane security configuration using the initial setup wizard. These settings can be changed at any time using the WebUI or the command-line interfaces.



If you are configuring control plane security for the first time after upgrading from ArubaOS 5.0 or earlier, see Configuring Control Plane Security after Upgrading on page 127 for details on enabling this feature using the WebUI or CLI.

In the WebUI

1. Access the WebUI of a standalone or master controller, and navigate to **Configuration > Network > Controller**.
2. Select the **Control Plane Security** tab.
3. Configure the following control plane security parameters:

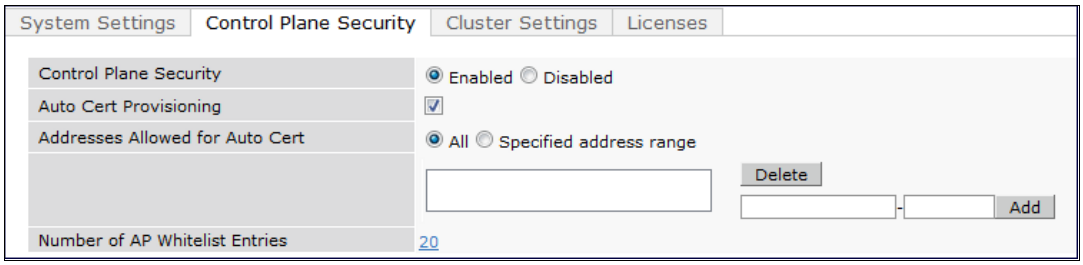
Table 12: *Control Plane Security Parameters*

Parameter	Description
Control Plane Security	Select **enable** or **disable** to turn the control plane security feature on or off. This feature is enabled by default.
Auto Cert Provisioning	When you enable the control plane security feature, you can select this checkbox to turn on automatic certificate provisioning. When you enable this feature, the controller attempts to send certificates to all associated campus APs. Auto certificate provisioning is disabled by default. **NOTE:** If you do not want to enable automatic certificate provisioning the first time you enable control plane security on the controller, you must identify the valid APs on your network by adding those to the campus AP whitelist. For details, see Viewing and Managing the Master or Local Controller Whitelists on page 118. After you have enabled automatic certificate provisioning, you must select either **Auto Cert Allow all** or **Addresses Allowed for Auto Cert**.
Addresses allowed for Auto Cert	The Addresses Allowed for Auto Cert section allows you to specify whether certificates are sent to all associated APs, or just APs within one or more specific IP address ranges. If your controller has a publicly accessible interface, you should identify your campus and Remote APs by IP address range. This prevents the controller from sending certificates to external or rogue campus APs that may attempt to access your controller through that interface. Select **All** to allow all associated campus and remote APs to receive automatic certificate provisioning. This parameter is enabled by default. Select **Addresses Allowed for Auto Cert** to send certificates to a group of campus or remote APs within a range of IP addresses. In the two fields below, enter the start and end IP addresses, then click **Add**. Repeat this procedure to add additional IP ranges to the list of allowed addresses. If you enable both control plane security and auto certificate provisioning, all APs in the address list receives automatic certificate provisioning. Remove a range of IP addresses from the list of allowed addresses by selecting the IP address range from the list and clicking **Delete.**
Number of AP Whitelist Entries	This parameter is the total number of APs in the remote AP and campus AP Whitelists. This number is also a link to a combined whitelist that displays all campus and remote AP entries.

4. Click **Apply** .

The master controller generates its self-signed certificate and begins distributing certificates to campus APs and any local controllers on the network over a clear channel. After all APs have received a certificate and have connected to the network using a secure channel, access the **Control Plane Security** window and turn off auto certificate provisioning if that feature was enabled. This prevents the controller from issuing a certificate to any rogue APs that may appear on your network at a later time.

Figure 4 *Control Plane Security Settings*



In the CLI

Use the commands below to configure control plane security via the command line interface on a standalone or master controller. Descriptions of the individual parameters are listed in Table 12, above.

```
control-plane-security
   auto-cert-allow-all
   auto-cert-allowed-addrs <ipaddress-start> <ipaddress-end>
   auto-cert-prov
   cpsec-enable
```

Example:

```
(host)(config) # control-plane-security
   auto-cert-prov
   no auto-cert-allow-all
   auto-cert-allowed-addrs 10.21.18.10 10.21.10.90
```

View the current control plane security settings using the following command:

```
   show control-plane-security
```

Managing AP Whitelists

Campus and Remote APs appear as valid APs in the campus and Remote AP whitelists when you manually enter their information into the whitelists via the controller's CLI or WebUI, or after the controller sends the AP a certificate via automatic certificate provisioning, and the AP connects to its controller via a secure tunnel. Any APs not approved or certified on the network are also included in the whitelists, but these APs appear in an unapproved state.

Use the whitelists to grant valid APs secure access to the network, or to revoke access from suspected rogue APs. When you revoke or remove an AP from the campus or remote AP whitelist on a controller that uses control plane security, that AP is not able to communicate with the controller again, except to obtain a new certificate.



If you manually add APs to the whitelists (rather than automatically adding the APs via the automatic certificate provisioning feature), make sure that the whitelists have been synchronized to all other controllers on the network *before* enabling control plane security.

Adding APs to the Campus and Remote AP Whitelists

You can add an AP to the campus AP or remote AP whitelists via the WebUI or command-line interface. To add an entry via the WebUI, use the following procedure:

1. Access the WebUI, and navigate to **Configuration > Wireless > AP Installation**.
2. Click the **Whitelist** tab.
3. Select the whitelist to which you want to add the AP. The Whitelist tab displays status information for the Campus AP Whitelist by default. To add a remote AP to the Remote AP whitelist, click the blue **Remote AP** link at the top of the table before you proceed to step 4 on page 111.

Figure 5 *Control Plane Security Settings*



4. Click **Entries** in the upper right corner of the whitelist status window.
5. Click **New**.
6. Define the following parameters for each AP you want to add to the whitelist.

Table 13: *AP Whitelist Parameters*

Parameter	Description
Campus AP whitelist configuration parameters	
AP MAC Address	MAC address of a campus AP that supports secure communications to and from its controller.
Description	(Optional) A brief description of the campus AP.
Remote AP whitelist configuration parameters	
AP MAC Address	MAC address of the remote AP, in colon-separated octets.
User Name	Name of the end user who provisions and uses the remote AP.
AP Group	Name of the AP group to which the remote AP is assigned.
AP Name	Name of the remote AP. If you not specify a name, the AP uses its MAC address as a name (Optional).
Description	A brief description to help you identify the AP (Optional).
IP-Address	The static inner IP address to be assigned to the remote APs.

7. Click **Add** .
8. Click **Apply**.

To add an AP to the Campus AP whitelist via the command-line interface, issue the command

```
whitelist-db cpsec add mac-address <macaddr> description <description>
```

To add an AP to the Remote AP whitelist via the command-line interface, issue the command

```
whitelist-db rap add mac-address <macaddr> ap-group <ap-group> [ap-name <ap-name>]
[description <description>] [full-name <name>] remote-ip <inner-ip-adr>
```

Viewing Whitelist Status

The WebUI can display either a table of entries in the selected whitelist, or a general nstatus summary for that whitelist. The whitelist status pages show the current status each entry in the whitelist, and, for controllers in a master/local controller topology, information for whitelist synchronization between controllers. This information is updated automatically as the status of each entry changes.

The **Wireless > AP Installation > Whitelist** tab displays status information for the campus AP Whitelist by default. To view status information for entries in the remote AP whitelist, click the blue **Remote AP** link on this tab.

The following table describes the status information types available on the Whitelist status page.

Table 14: *Whitelist status information*

Status Entry	Description
Control Plane Security (Campus AP Whitelist status only)	Shows if control plane security has been enabled or disabled on the controller. This status entry is also a link to the control plane security configuration tab.
Number of Entries	Total number of entries in the selected whitelist.
Approved Entries	Number of entries that have been approved by the controller.
Unapproved Entries	Number of entries that have not been approved by the controller
Certified Entries	AP has an approved certificate from the controller.
Certified Hold Entries	Shows if the controller thinks the AP has been certified with a factory certificate yet the AP requests to be certified again. Because this is not a normal condition, the AP is not approved as secure until a network administrator manually changes the status of the AP to verify that it is not compromised. **NOTE:** If an AP is in this state due to connectivity problems, then the AP recovers and is taken out of this hold state as soon as connectivity is restored.
Revoked Entries	Number of AP entries that have been manually revoked.
Marked For Deletion Entries	Number of APs that have been marked for deletion, but that have not been removed from the whitelist.

The Remote AP whitelist entries page displays only the information you manually configure. The entries in the campus AP whitelist include both user-defined settings and additional AP information that is updated as the status of the AP changes.

Table 15: *Additional Campus AP Status Information*

Parameter	Description
Cert Type	The type of certificate used by the AP. • **switch-cert**: The AP is using a certificate signed by the controller. • **factory-cert**: The AP is using a factory-installed certificate.
State	The Campus AP Whitelist reports one of the following states for each campus AP: • **unapproved-no-cert:** The AP has no certificate and is not approved. • **unapproved-factory-cert:** The AP has a preinstalled certificate that was not approved. • **approved-ready-for-cert:** The AP has been approved as a valid campus AP and is ready to receive a certificate. • **certified-factory-cert**: The AP is already has a factory certificate. If an AP has the **factory-cert** certificate type and is in the **certified-factory-cert** state, then that campus AP is not reissued a new certificate if you enable automatic certificate provisioning. • **certified-switch-cert:** The AP has an approved certificate from the controller. • **certified-hold-factory-cert**: An AP is put in this state when the controller thinks the AP has been certified with a factory certificate but the AP requests to be certified again. Because this is not a normal condition, the AP is not approved as a secure AP until a network administrator manually changes the status of the AP to verify that it is not compromised. **NOTE:** If an AP is in this state due to connectivity problems, then the AP recovers and leaves this hold state as soon as connectivity is restored. • **certified-hold-switch-cert**: An AP is put in this state when the controller thinks the AP has been certified with a controller certificate but the AP requests to be certified again. Because this is not a normal condition, the AP is not approved as a secure AP until a network administrator manually changes the status of the AP to verify that it is not compromised. **NOTE:** If an AP is in this state due to connectivity problems, then the AP recovers and is taken out of this hold state as soon as connectivity is restored.
Revoked	Shows if the AP's secure status has been revoked.
Revoked Text	An optional, brief statement describing why the AP was revoked.
Last Update	Time and date of the last AP status update.

To view information about the remote and campus AP whitelists using the command-line interface, use the commands described in Table 16.

Table 16: *View the Campus AP Whitelist via the CLI*

Command	Description
`show whitelist-db cpsec [mac-address <macaddr>]`	Shows detailed information for each AP in the whitelist, including the AP's MAC address, approved state, certificate type, and description. Include the optional **`mac-address <macaddr>`** parameters to view data for a single entry.
`show whitelist-db cpsec-status`	The command gives aggregate information for the numbers of APs in each of the following categories: • Total entries • Approved entries • Unapproved entries • Certified entries • Certified hold entries • Revoked entries • Marked for deletion entries

Modifying an AP in the Campus AP Whitelist

Use the following procedure to modify a campus AP entry's certificate type, state, description, and revoked status via the WebUI:

1. Access the master controller WebUI, and navigate to **Configuration>AP Installation**.
2. Click the **Campus AP Whitelist** tab.
3. Select the checkbox by the entry for the AP you want to edit, then click **Modify**.

   If your campus AP whitelist is large and you cannot immediately locate the AP entry you want to edit, select the **Search** link by the upper right corner of the whitelist. The **Campus AP Whitelist** tab displays several fields that allow you to search for an AP with a specified MAC address, certificate type or state. Specify the values that match the AP you want locate, then click **Search** . The whitelist displays a list of APs that match your search criteria. Select the AP from this list, then click **Modify**.
4. Update the AP's whitelist entry with the new settings. Some of the configurable parameters were available when you first defined the entry, and are described in Table 13 above. When you modify an existing whitelist entry, you can also configure the following additional parameters that were not configurable when you first created the entry:
   - **Cert-type**: The type of certificate used by the AP.
     - switch-cert: The campus AP is using a certificate signed by the controller.
     - factory-cert: The campus AP is using a factory-installed certificate.
   - **State**: When you click the **State** drop-down list to modify this parameter, you may choose one of the following options:
     - approved-ready-for-cert: The AP has been approved state and is ready to receive a certificate.
     - certified-factory-cert: The AP is certified and has a factory-installed certificate.
   - **Revoke**: Click the **Revoke** checkbox to revoke an AP's secure status. When you select this checkbox, you can enter a brief comment explaining why the AP is being revoked.
5. Click **Update** to update the campus AP whitelist entry with its new settings.

To modify an entry in the campus AP whitelist via the command-line interface, issue the following commands:

```
whitelist-db cpsec modify mac-address
   cert-type switch-cert|factory-cert
   description <description>
   mode disable|enable
   revoke-text <revoke-text>
```

```
state approved-ready-for-cert|certified-factory-cert
```

Revoking an AP via the Campus AP Whitelist

You can revoke an invalid or rogue AP either by opening the **modify** menu and modifying the AP's revoke status (as described in the section above), or by selecting the AP in the campus whitelist and revoking its secure status directly, without modifying any other parameters or entering a description of why that AP was revoked. When you revoke an AP's secure status in the campus AP whitelist, the whitelist retains the AP's status information. To revoke an invalid or rogue AP and permanently remove the AP from the whitelist, you must delete that entry.

To revoke an AP via the WebUI:

1. Access the master controller WebUI, and navigate to **Configuration > AP Installation**.
2. Click the **Campus AP Whitelist** tab.
3. To revoke one or more secure campus APs, select the checkbox by the entry for each AP whose secure status should be revoked, then click **Revoke**.

   If your campus AP whitelist is large and you cannot immediately locate the AP entry you want to revoke, select the **Search** link by the upper right corner of the whitelist. The **Campus AP Whitelist** tab displays several fields that allow you to search for an AP with a specified MAC address, certificate type, or state. Specify the values that match the AP you want locate, then click **Search** . The whitelist displays a list of APs that match your search criteria. Select the AP from this list, then click **Revoke**.

To revoke an AP via the command-line interface, issue the command:

```
whitelist-db cpsec revoke mac-address <macaddr> revoke-text <"revoke text">
```

Deleting an AP Entry from the Campus AP Whitelist

Before you delete an AP entry from the campus whitelist, verify that auto certificate provisioning is either no longer enabled, or only enabled for IP addresses that do not include the AP being removed. If you enable automatic certificate provisioning for an AP that it is still connected to the network, you cannot permanently delete it from the campus AP whitelist; the controller immediately recertifies the AP and recreates its whitelist entry.

To delete an AP entry via the WebUI:

1. Access the master controller WebUI, and navigate to **Configuration > AP Installation**.
2. Click the **Campus AP Whitelist** tab.
3. Select the checkbox by entry for each AP you want to remove, then click **delete**.

   If your campus AP whitelist is large and you cannot immediately locate the AP entry you want to delete, select the **Search** link by the upper right corner of the whitelist. The **Campus AP Whitelist** tab displays several fields that allow you to search for an AP with a specified MAC address, certificate type, or state. Specify the values that match the AP you want locate, then click **Search** . The whitelist displays a list of APs that match your search criteria. Select the AP from this list, then click **delete**.

To delete an AP entry via the CLI, issue the command:

```
whitelist-db cpsec del mac-address <macaddr>
```

Purging the Campus AP Whitelist

Before you add a new local controller to a network using control plane security, you must purge the campus AP whitelist on the new controller. As soon as you add the new controller to the hierarchy, the entries in the new controller's campus AP whitelist merge into the whitelist for all other master and local controllers. If you add any old or invalid AP entries to the campus AP whitelist, all controllers in the hierarchy will trust those APs,

creating a potential security risk. For additional information on adding a new local controller using control plane security to your network, see Replacing a Local Controller on page 123

To purge a controller's campus AP whitelist via the WebUI:

1. Access the master controller WebUI, and navigate to **Configuration > AP Installation**.
2. Click the **Campus AP Whitelist** tab.
3. Click **Purge**.

To purge a campus AP whitelist via the command-line interface, issue the command:

```
whitelist-db cpsec purge
```

OffLoading a Controller RAP Whitelist to ClearPass Policy Manager

This feature allows whitelist entries for remote APs (RAPs) to be maintained externally in a ClearPass Policy Manager (CPPM) server. The controller, if configured to use an external server, can send a RADIUS access request to a CPPM server. The RAP MAC address is used as a username and password to construct the access request packet and the CPPM validates the RADIUS message and returns the relevant parameters for the authorized RAPs.

The following three supported parameters are associated with the following VSAs. They are sent by the CPPM server in the RADIUS access accept packet for authorized RAPs:

- ap-group: Dell-AP-Group
- ap-name: Dell-Location-ID
- remote-ip: Dell-AP-IP-Address

The following defaults are used when any of the supported parameters are not provided by the CPPM server in the RADIUS access accept response:

- ap-group: The default ap-group is assigned to the RAP.
- ap-name: The RAP MAC address is used as the AP name.
- remote-ip: The controller selects the remote IP address from its available pool of addresses.

There is no change in the RAP role assignment. The RAP is assigned the role that is configured in the VPN *default-rap* profile.

In the WebUI

To assign a CPPM server to a RAP:

1. Configure a CPPM server using the controller WebUI:
   a. Navigate to the **Configuration > Security > Authentication > Servers page**.
   b. Select **Radius Server** to display the CPPM Server List.
   c. To configure a CPPM server, enter the name for the server and click **Add**.
   d. Select the name to configure server parameters. Select the **Mode** check box to activate the authentication server.
   e. Click **Apply**.
2. Create a server group that contains the CPPM server.
3. Navigate to **Configuration > All Profile Managment > Wireless LAN > VPN Authentication > default-rap > Server Group**.
4. Select the CPPM server from the Server Group drop-down list.
5. Click **Apply**.

To assign a CPPM server to a RAP that was initially an IAP:

1. Make sure that a CPPM server is configured on the controller.
2. Navigate to **Configuration > All Profile Managment > Wireless LAN > VPN Authentication > default-iap > Server Group**.
3. Select the CPPM server from the Server Group drop-down list.
4. Click **Apply**.

In the CLI

Configure a radius server with CPPM server as host address. In this example **cppm-rad** is the CPPM server name and **cppm-sg** is the server group name.

```
(host)(config) #aaa authentication-server radius cppm-rad
```

Add this server to a server group:

```
(host)(config) #aaa server-group cppm-sg
  auth-server cppm-rad
```

Add this server group to the **default-rap** vpn profile:

```
(host)(config) #aaa authentication vpn default-rap
  server-group cppm-sg
```

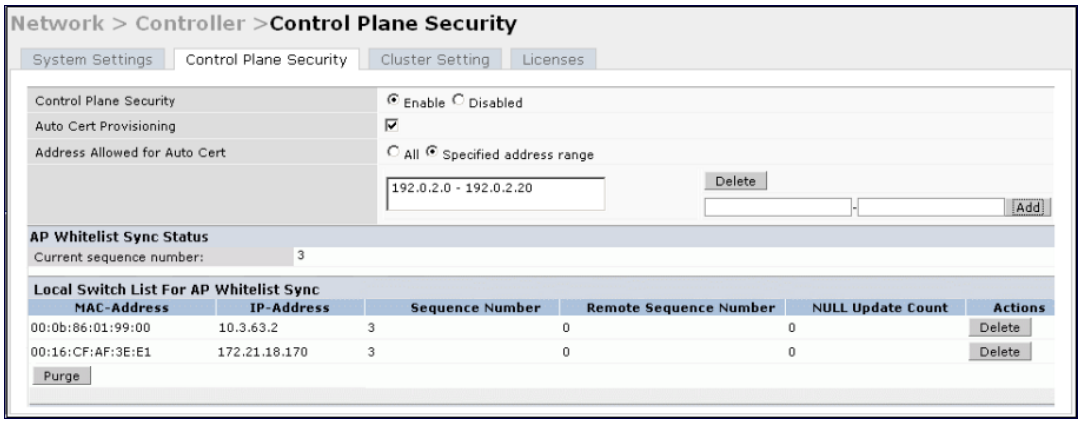
Managing Whitelists on Master and Local Controllers

Every controller using the control plane security feature maintains a campus AP whitelist, a local controller whitelist and a master controller whitelist. The contents of these whitelists vary, depending upon the role of the controller, as shown in the figure below.

Table 17: *Control Plane Security Whitelists*

Controller Role	Campus AP Whitelist	Master Controller Whitelist	Local Controller Whitelist
On a (standalone) master controller with no local controllers:	The campus AP whitelist contains entries for the secure campus APs associated with that controller.	The master controller whitelist is empty, and does not appear in the WebUI.	The local controller whitelist is empty, and does not appear in the WebUI.
On a master controller with local controllers:	The campus AP whitelist contains an entry for every secure campus AP on the network, regardless of the controller to which it is connected.	The master controller whitelist is empty, and does not appear in the WebUI.	The local controller whitelist contains an entry for each associated local controller.
On a local controller:	The campus AP whitelist contains an entry for every secure campus AP on the network, regardless of the controller to which it is connected.	The master controller whitelist contains the MAC and the IP addresses of the master controller.	The local controller whitelist is empty, and does not appear in the WebUI.

Figure 6 *Local Controller Whitelist on a Master Controller*



If your deployment includes both master and local controllers, then the campus AP whitelist on every controller contains an entry for every secure AP on the network, regardless of the controller to which it is connected. The master controller also maintains a whitelist of local controllers using control plane security. When you change a campus AP whitelist on any controller, that controller contacts the other connected controllers to notify them of the change.

The master controller whitelist on each local controller contains the IP and MAC addresses of its master controller. If your network has a redundant master controller, then this whitelist contains more than one entry. You rarely need to delete the master controller whitelist. Although you can delete an entry from the master controller whitelist, you should do so only if you have removed a master controller from the network.

Campus AP Whitelist Synchronization

The current sequence number in the **AP Whitelist Sync Status** field shows the number of changes to the campus AP whitelist made on that controller. Each controller compares its campus AP whitelist against whitelists on other controllers every two minutes by default. If a controller detects a difference, it sends its changes to the other controllers on the network. If all other controllers on the network have successfully received and acknowledged all whitelist changes made on that controller, every entry in the **sequencenumber** column in the local controller or master controller whitelists has the same value as the sequence number displayed in the **AP Whitelist Sync Status** field. If a controller in the master or local controller whitelist has a lower sequence number, that controller may still be waiting to complete its update, or receive its update acknowledgement. In the example in Figure 6, the master controller has a current sequence number of 3, and each sequence number in its local controller whitelist also shows a value of 3, indicating that both local controllers have received and acknowledged all three campus AP whitelist changes made on the master controller. For additional information on troubleshooting whitelist synchronization, see Verifying Whitelist Synchronization on page 129.

You can view a controller's current sequence number via the CLI using the command:

```
show whitelist-db cpsec-seq
```

Viewing and Managing the Master or Local Controller Whitelists

The following sections describe the commands to view and delete entries in a master or local controller whitelist.

Viewing the Master or Local Controller Whitelist

To view the master or local controller whitelists via the WebUI, use the procedure below:

1. Access the controller's WebUI, and navigate to **Configuration > AP Instalation.**
2. Select the **Whitelist** tab.

The master and local controller tables each include the following information:

Table 18: *Master and Local Controller Whitelist Information*

Data Column	Description
MAC-Address	On a local controller whitelist: MAC address of the master controller. On a master controller whitelist: MAC address of a local controller.
IP-Address	On a local controller whitelist: IP address of the master controller. On a master controller whitelist: IP address of a local controller.
Sequence Number	The number of times the controller in the whitelist received and acknowledged a campus AP whitelist change from the controller whose WebUI you are currently viewing. For deployments with both master and local controllers: • The sequence number on a master controller should be the same as the remote sequence number on the local controller. • The sequence number on a local controller should be the same as the remote sequence number on the master controller.
Remote Sequence Number	The number of times that the controller whose WebUI you are viewing received and acknowledged a campus AP whitelist change from the controller in the whitelist. For deployments with both master and local controllers: • The remote sequence number on a master controller should be the same as the sequence number on the local controller. • The remote sequence number on a local controller should be the same as the sequence number on the master controller.
Null Update Count	The number of times the controller checked its campus AP whitelist and found nothing to synchronize with the other controller. The controller compares its control plane security whitelist against whitelists on other controllers every two minutes by default. If the null update count reaches five, the controller sends an "empty sync" heartbeat to the remote controller to ensure the sequence numbers on both controllers are the same, then resets the null update count to zero.

To view the master or local controller whitelists via the command-line interface, issue the following commands:

```
show whitelist-db cpsec-master-switch-list [mac-address <mac-address>]
show whitelist-db cpsec-local-switch-list [mac-address <mac-address>]
```

Deleting an Entry from the Master or Local Controller Whitelist

You do not need to delete a master controller from the master controller whitelist during the course of normal operation. However, if you remove a local controller from the network, you should also remove the local controller from the local controller whitelist on the master controller. If the local controller whitelist contains entries for controllers no longer on the network, then a campus AP whitelist entry can be marked for deletion but is not physically deleted, as the controller is waiting for an acknowledgment from another controller no longer on the network. This can increase network traffic and reduce memory resources on the controller.

To delete an entry from the master or local controller whitelist via the WebUI:

1. Access the controller's WebUI, and navigate to **Configuration > Controller.**
2. Select the **Control Plane Security** tab.
3. To delete an entry from the Local Controller Whitelist: In the **Local Switch List For AP Whitelist Sync** section, click the **Delete** button by each controller entry you want to remove.

   Or,

To delete an entry from the Master Controller Whitelist: In the **Master Switch List For AP Whitelist Sync** section, click **Delete** by each controller entry you want to remove.

4. Click **Apply**.

To delete an entry from the master or local controller whitelist via the command-line interface, issue either of the following commands:

```
whitelist-db cpsec-master-switch-list del mac-address <mac-address>
whitelist-db cpsec-local-switch-list del mac-address <mac-address>
```

Purging the Master or Local Controller Whitelist

There is no need to purge a master controller whitelist during the course of normal operation. If, however, you are removing a controller from the network, you can purge its controller whitelist after it has been disconnected from the network. To clear a local controller whitelist entry on a master controller that is still connected to the network, select that individual whitelist entry and delete it using the **delete** option.

To purge a controller whitelist via the WebUI, use the following procedure:

1. Access the controller's WebUI, and navigate to **Configuration > Controller**.
2. Select the **Control Plane Security** tab.
3. To clear the Local Controller Whitelist: In the **Local Switch List For AP Whitelist Sync** section, click **Purge.**

   Or,
4. To clear the Master Controller Whitelist: In the **Master Switch List For AP Whitelist Sync** section, click **Purge**.

To purge a controller whitelist via the command-line interface, issue the following commands:

```
whitelist-db cpsec-master-switch-list purge
whitelist-db cpsec-local-switch-list purge
```

Working in Environments with Multiple Master Controllers

Configuring Networks with a Backup Master Controller

If your network includes a redundant backup master controller, you *must synchronize the database from the primary master to the backup master at least once* after all APs are communicating with their controllers over a secure channel. This ensures that all certificates, IPsec keys, and campus AP whitelist entries are synchronized to the backup controller. You should also synchronize the database any time the campus AP whitelist changes (APs are added or removed to ensure that the backup controller has the latest settings).

Master and backup controllers can be synchronized using either of the following methods:

- **Manual Synchronization**: Issue the **database synchronize** CLI command in enable mode to manually synchronize databases from your primary controller to the backup controller.
- **Automatic Synchronization**: Schedule automatic database backups using the **database synchronize period** CLI command in config mode.

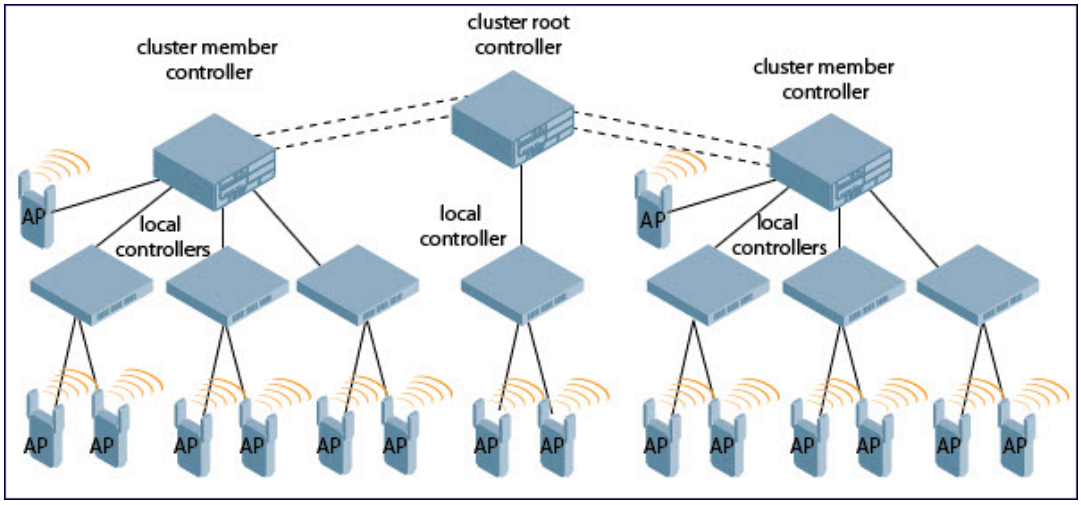


If you add a new backup controller to an existing controller, you must add the backup controller as the **lower priority** controller. If you do not add the backup controller as a lower priority controller, your control plane security keys and certificates may be lost. If you want the new backup controller to become your primary controller, increase the priority of that controller to a primary controller *after* you have synchronized your data.

Configuring Networks with Clusters of Master Controllers

If your network includes multiple master controllers each with their own hierarchy of APs and local controllers, you can allow APs from one hierarchy to failover to any other hierarchy by defining a *cluster* of master controllers. Each cluster has one master controller as its cluster root, and all other master controllers as cluster members. The master controller operating as the cluster root creates a self-signed certificate, then certifies its own local controllers and APs. Next, the cluster root sends a certificate to each cluster member, which in turn certifies its own local controllers and APs. Because all controllers and APs in the cluster have the same trust anchor, the APs can switch to any other controller in the cluster and still remain securely connected to the network.

Figure 7 *A Cluster of Master Controllers using Control Plane Security*



To create a controller cluster, you must first define the root master controller and set an IPsec key or select a certificate for communications between the cluster root and cluster members.

NOTE

You must use the command-line interface to configure certificate authentication for cluster members. The WebUI supports cluster authentication using IPsec keys only. If your master and local controllers use a pre-shared key for authentication, they create the IPsec tunnel using IKEv1. If your master and local controllers use certificates for authentication, the IPsec tunnel is created using IKEv2.

Creating a Cluster Root

Use the WebUI to identify a controller as a cluster root, and use an IPsec key to secure communication between the cluster root and cluster members. Use the command-line interface to create a cluster root using an IPsec key, factory-installed certificate, or custom certificate.

To create a cluster root using the WebUI:

1. Access the WebUI of the controller you want to identify as the cluster root, and navigate to **Configuration > Controller**.
2. Click the **Cluster Setting** tab.
3. For the cluster role, select **Root.**
4. In the **Cluster Member IPsec Keys** section, enter the controller IP address of a member controller in the cluster. If you want to use a single key for all member controllers, use the IP address **0.0.0.0**.
5. In the **IPsec Key** and **Retype IPsec Key** fields, enter the IPsec key for communication between the specified member controller and the cluster root.
6. Click **Add**.
7. *Optional*: repeat steps 4-6 to add another member controller to the cluster.

8. Click **Apply.**

To create a cluster root via the CLI, access the command-line interface of the controller you want to identify as the root of the controller cluster, then issue one of the following commands:

- To authenticate cluster members using a custom certificate:

  ```
  cluster-member-custom-cert member-mac <mac> ca-cert <ca> server-cert <cert>     suite-b <gcm-128 | gcm-256>]
  ```

- To authenticate cluster members using a factory-installed certificate:

  ```
  cluster-member-factory-cert member-mac <mac>
  ```

- To authenticate cluster members using an IPsec key:

  ```
  cluster-member-ip <ip-address> ipsec <key>
  ```

  The `<ip-address>` parameter in this command is the IP address of a member controller in the cluster, and the **`<key>`** parameter in each command is the IPsec key for communication between the specified member controller and the cluster root. Use the IP address **0.0.0.0** in this command to set a single IPsec key for all member controllers, or repeat this command as desired to define a different IPsec key for each cluster member.

Creating a Cluster Member

Once you have identified the cluster root, you must then identify the member controllers in the cluster.

Use the WebUI to identify a controller as a cluster member, and use an IPsec key to secure communication between the cluster member and the cluster root. Use the command-line interface to create a cluster member and secure communications between that member and the cluster root using an IPsec key, factory-installed certificate, or custom certificate.

To create a cluster member using the WebUI:

1. Access the WebUI of the cluster member controller, and navigate to **Configuration > Controller**.
2. Click the **Cluster Setting** tab.
3. For the cluster role, select **Member**.
4. In the **Controller IP Address** field, enter the IP address of the root controller in the cluster.
5. In the **IPsec Key** and **Retype IPsec Key** fields, enter the IPsec key for communication between the specified member controller and the cluster root. This parameter must be have the same value as the key defined for the cluster member in Creating a Cluster Root on page 121.
6. Click **Add**.
7. Click **Apply**.

To create a cluster root via the CLI, access each of the member master controllers and define the IPsec key or certificate for communication between that controller and the cluster root.

```
cluster-root-ip <ip-address>
   ipsec <key>
   factory-cert master-mac <mac>
   ipsec-custom-cert master-mac1 <mac1> [master-mac2 <mac2>] ca-cert <ca> server-cert <cert>
   [suite-b <gcm-128 | gcm-256>]
```

In this command the **<ip-address>** parameter is the IP address of the root master controller in the cluster. If you are using an IPsec key, the **<key>** parameter in this command must be have the same value as the key defined for the cluster member via the **cluster-member-ip** command.

Viewing Controller Cluster Settings

To view your current cluster configuration via the WebUI:

1. Navigate to **Configuration > Controller**.
2. Click the **Cluster Setting** tab.

- If you are viewing the WebUI of a cluster root, the output of this command displays the IP address of the VLAN on the cluster member used to connect to the cluster root.
- If you are viewing the WebUI of a cluster member, the output of this command displays the IP address of the VLAN on the cluster root used to connect to the cluster member.

To view your current cluster configuration via the command-line interface, issue the CLI commands described in Table 19.

Table 19: *CLI Commands to Display Cluster Settings*

Command	Description
show cluster-switches	When you issue this command from the cluster *root*, the output of this command displays the IP address of the VLAN the cluster member uses to connect to the cluster root. If you issue this command from a cluster *member*, the output of this command displays the IP address of the VLAN the cluster root uses to connect to the cluster member.
show cluster-config	When you issue this command from the cluster *root*, the output of this command shows the cluster role of the controller, and the IP address of each active member controller in the cluster. When you issue this command from a cluster *member*, the output of this command shows the cluster role of the controller, and the IP address of the cluster root.

Replacing a Controller on a Multi-Controller Network

The procedure to replace a controller within a multi-controller network varies, depending upon the role of that controller, whether the network has a single master controller or a cluster of master controllers, and whether or not the controller has a backup.



The following sections describe the steps to replace an existing controller. To add a new local controller to a network, or to permanently remove a local controller without replacing it, see Viewing and Managing the Master or Local Controller Whitelists on page 118.

Replacing Controllers in a Single Master Network

Use the procedures in this section to replace a master or local controller in a network environment with a single master controller.

Replacing a Local Controller

Use the following procedure to replace a local controller in a single-master network:

1. Disconnect the local controller from the network.
2. If you plan on moving the local controller to another location on the network, purge the campus AP whitelist on the controller.

   Access the command-line interface on the old local controller and issue the command **`whitelist-db cpsec purge`**
   or,
   Access the local controller WebUI, navigate to **Configuration > AP Installation > Campus AP Whitelist** and click **Purge**.
3. Once you purge the campus AP whitelist, you must inform the master controller that the local controller is no longer available using one of these two methods:

This step is very important; unused local controller entries in the local controller whitelist can significantly increase network traffic and reduce controller memory resources.

- Access the command-line interface on the master controller, and issue the command **whitelist-db cpsec-local-switch-list del mac-address <local-controller-mac>**
- Access the master controller WebUI, navigate to the **Configuration > Controller > Control Plane Security** window, select the entry for the local controller you want to delete from the local controller whitelist, and click **Delete**.

4. Install the new local controller, but do not connect it to the network yet. If the controller has been previously installed on the network, you must ensure that the new local controller has a clean whitelist.
5. Purge the local controller whitelist using one of the following two methods:
   - Access the command-line interface on the new local controller and issue the command **whitelist-db cpsec purge**
   - Access the local controller WebUI, navigate to **Configuration > AP Installation > Campus AP Whitelist** and click **Purge**.
6. Now connect the new local controller to the network. It is very important that the local controller be able to contact the master controller the first time it connects to the network, because the master controller certifies the local controller's control plane security certificate the first time the local controller contacts its master.
7. Once the local controller has a valid control plane security certificate and configuration, the local controller receives the campus AP whitelist from the master controller and starts certifying approved APs.
8. APs associated with the new local controller reboots and creates new IPsec tunnels to their controller using the new certificate keys.

Replacing a Master Controller with No Backup

Use the following procedure to replace a master controller that does not have a backup controller:

1. Remove the old master controller from the network.
2. Install and configure the new master controller, then connect the new master to the network. The new master controller generates a new certificate when it first becomes active.
3. If the new master controller has a different IP address than the old master controller, change the master IP address on the local controllers to reflect the address of the new master.
4. Reboot each local controller to ensure the local controllers obtain their certificate from the new master. Each local controller begins using a new certificate signed by the master controller.
5. APs are now no longer able to securely communicate with the controller using their current key, and must obtain a new certificate. Access the campus AP whitelist on any local controller, and change all APs in a “certified” state to an “approved” state. The new master controller sends the approved APs new certificates. The APs reboot and create new IPsec tunnels to their controller using the new certificate key.

   If the master controller does not have any local controllers, you must recreate the campus AP whitelist by turning on automatic certificate provisioning or manually reentering the campus AP whitelist entries.

Replacing a Redundant Master Controller

The control plane security feature requires you to synchronize databases from the primary master controller to the backup master controller at least once after the network is up and running. This ensures that all certificates, keys, and whitelist entries are synchronized to the backup controller. Because the AP whitelist may change periodically, you should regularly synchronize these settings to the backup controller. For details, see Configuring Networks with a Backup Master Controller on page 120.

When you install a new backup master controller, *you must add it as a lower priority* controller than the existing primary controller. After you install the backup controller on the network, synchronize the database from the existing primary controller to the new backup controller to ensure that all certificates, keys, and whitelist entries required for control plane security are added to the new backup controller configuration. If you want the new controller to act as the primary controller, you can increase that controller's priority *after* the settings have been synchronized.

Replacing Controllers in a Multi-Master Network

Use the following procedures to replace a master or local controller in a network environment with a multiple master controllers.

Replacing a Local Controller in a Multi-Master Network

The procedure to replace a local controller in a network with multiple master controllers is the same as the procedure to replace a local controller in a single-master network. To replace a local controller in a multi-master network, follow the procedure described in Replacing a Local Controller on page 123

Replacing a Cluster Member Controller with no Backup

The control plane security feature allows APs to fail over from one controller to another within a cluster. Therefore, cluster members or their local controllers may have associated APs that were first certified under some other cluster member (or the cluster root). If you permanently remove a cluster member whose APs were all originally certified under the cluster member being removed, its associated APs do not need to reboot in order to connect to a different controller. If, however, you remove a cluster member whose associated APs were originally certified under a *different* cluster member, those APs need to reboot and be recertified before they can connect to a different controller. If the cluster member you are removing has local controllers, the local controllers also reboot so they can be updated with new certificates, then pass the trust update to their terminating APs.

To replace a cluster member that does not have a backup controller:

1. On the cluster master to be removed, clear the cluster root IP address by accessing the command-line interface and issuing the command **no cluster-root-ip <cluster-root-ip> ipsec <clusterkey>**.
2. Remove the cluster member from the network.
3. If the cluster master you removed has any associated APs, you must reboot those APs so they receive an updated certificate.
4. If the cluster member you removed has any associated local controllers, reboot those local controllers so they receive a new certificate and then pass that trust update to their APs.
5. Remove the cluster master from the cluster root's master controller list by accessing the command-line interface on the cluster root and issuing the command **whitelist-db cpsec-master-switch-list del mac-address <cluster-master-mac>**.



This step is very important. Unused local controller entries in the local controller whitelist can significantly increase network traffic and reduce controller memory resources.

6. Remove the old cluster member from the network. Remember, that controller still has campus AP whitelist entries from the entire cluster. You may want to delete or revoke unwanted entries from the campus AP whitelist.

Now, you must install the new cluster member controller according to the procedure described in Creating a Cluster Member on page 122. The new cluster member obtains a certificate from the cluster root when it first becomes active.

7. If the new cluster member has any associated APs, reboot those APs so they obtain a trust update.

8. If the new cluster member has any local controllers, reboot the local controllers associated with the new cluster member. The local controllers obtain a new certificate signed by the cluster member, and then pass that trust update to their associated APs.

Replacing a Redundant Cluster Member Controller

The control plane security feature requires you to synchronize databases from the primary controller to the backup controller at least once after the network is up and running. This ensures that all certificates, keys, and whitelist entries are synchronized to the backup controller. Because the AP whitelist may change periodically, you should regularly synchronize these settings to the backup controller. For details, see Configuring Networks with a Backup Master Controller on page 120.

When you install a new backup cluster member, *you must add it as a lower priority* controller than the existing primary controller. After you install the backup cluster member on the network, resynchronize the database from the existing primary controller to the new backup controller to ensure that all certificates, keys, and whitelist entries required for control plane security are added to the new backup controller configuration. If you want the new controller to act as the primary controller, you can increase that controller's priority *after* the settings have been resynchronized.

Replacing a Cluster Root Controller with no Backup Controller

If you replace a cluster root controller that does not have a backup controller, the new cluster root controller creates its own self-signed certificate. You then need to reboot each controller in the hierarchy in a specific order to certify all APs with that new certificate:

1. Remove the old cluster root from the network.
2. Install and configure the new cluster root.
3. Connect the new cluster root to the network so it can access cluster masters and local controllers.
4. If necessary, reconfigure the cluster masters and local controllers with their new cluster root IP and master IP addresses.
5. Reboot every cluster member controller. The cluster member begins using a new certificate signed by the cluster root.
6. Reboot every local controller. Each local controller begins using a new certificate signed by the cluster member.
7. Because the cluster root is new, it does not have a configured campus AP whitelist. Access the campus AP whitelist on any local controller or cluster master, and change all APs in a "certified" state to an "approved" state. The APs get recertified, reboot, and create new IPsec tunnels to their controller using the new certificate key.

   If a cluster root controller does not have any cluster master or local controllers, you must recreate the campus AP whitelist on the cluster root by turning on automatic certificate provisioning or manually reentering the campus AP whitelist entries.

Replacing a Redundant Cluster Root Controller

Best practices is to use a backup controller with your cluster root controller. If your cluster root has a backup controller, you can replace the backup cluster root without having to reboot all cluster master and local controllers, minimizing network disruptions.

The control plane security feature requires you to synchronize databases from the primary controller to the backup controller at least once after the network is up at running. This ensures that all certificates, keys, and whitelist entries are synchronized to the backup controller. Because the AP whitelist may change periodically, you should regularly synchronize these settings to the backup controller. For details, see Configuring Networks with a Backup Master Controller on page 120.

When you install a new backup cluster root, *you must add it as a lower priority* controller than the existing primary controller. After you install the backup cluster root on the network, resynchronize the database from the existing primary controller to the new backup controller to ensure that all certificates, keys, and whitelist entries required for control plane security are added to the new backup controller configuration. If you want the new controller to act as the primary controller, you can increase that controller's priority *after* the settings have been resynchronized.

Configuring Control Plane Security after Upgrading

When you initially deploy a controller running ArubaOS 6.0 or later, create your initial control plane security configuration using the initial setup wizard. However, if you are upgrading to ArubaOS 6.0 or if you are upgrading from ArubaOS 5.0 *but did not yet have control plane security enabled before the upgrade*, then you can use the strategies described in Table 20 to enable and configure control plane security feature.



If you upgrade a controller running ArubaOS 5.0.x to ArubaOS 6.0 or later, then the controller's control plane security settings do not change after the upgrade. If control plane security was already enabled, then it remains enabled after the upgrade. If it was not enabled previously, but you want to use the feature after upgrading, then you must manually enable it.

Table 20: *Control Plane Security Upgrade Strategies*

Automatically send Certificates to Campus APs	Manually Certify Campus APs
1. Access the control plane security window and enable both the control plane security feature and the auto certificate provisioning option. Next, specify whether you want all associated campus APs to automatically receive a certificate, or if you want to certify only those APs within a defined range of IP addresses.	1. Identify the campus APs that should receive certificates by entering the campus APs' MAC addresses in the campus AP whitelist.
2. Once all APs have received their certificates, disable auto certificate provisioning to prevent certificates from being issued to any rogue APs that may appear on your network at a later time.	2. If your network includes both master and local controllers, wait a few minutes, then verify that the campus AP whitelist has been propagated to all other controllers on the network. Access the WebUI of the master controller, navigate to **Configuration > Controller > Control Plane Security**, then verify that the **Current Sequence Number** field has the same value as the**Sequence Number** entry for each local controller in the local controller whitelist. (For details, see Verifying Whitelist Synchronization on page 129.)
3. If a valid AP did not receive a certificate during the initial certificate distribution, you can manually certify the AP by adding that AP's MAC address to the campus AP whitelist. You can also use this whitelist to revoke certificates from APs that should not be allowed access to the secure network.	3. Enable the control plane security feature.



If you upgraded your controller from ArubaOS 5.0 or earlier and you want to use this feature for the first time, you must either add all valid APs to the campus AP whitelist, or enable automatic certificate provisioning *before you enable the feature*. If you do not enable automatic certificate provisioning, only the APs currently approved in the campus AP whitelist are allowed to communicate with the controller over a secure channel. Any APs that do not receive a certificate will not be able to communicate with the controller except to request a certificate.

Troubleshooting Control Plane Security

Identifying Certificate Problems

If an AP has a problem with its certificate, check the state of the AP in the campus AP whitelist. If the AP is in either the certified-hold-factory-cert or certified-hold-switch-cert states, you may need to manually change the status of that AP before it can be certified.

- **certified-hold-factory-cert**: An AP is put in this state when the controller thinks the AP has been certified with a factory certificate, but the AP requests to be certified again. Because this is not a normal condition, the AP is not approved as a secure AP until you manually change the status of the AP to verify that it is not compromised. If an AP is in this state due to connectivity problems, then the AP recovers and is taken out of this hold state as soon as connectivity is restored.
- **certified-hold-switch-cert**: An AP is put in this state when the controller thinks the AP has been certified with a controller certificate yet the AP requests to be certified again. Because this is not a normal condition, the AP is not be approved as a secure AP until a network administrator manually changes the status of the AP to verify that it is not compromised. If an AP is in this state due to connectivity problems, then the AP recovers and is taken out of this hold state as soon as connectivity is restored.

Verifying Certificates

If you are unable to configure the control plane security feature on W-600 Series, W-6000M3, or W-3000 Series controllers, verify that its Trusted Platform Module (TPM) and factory-installed certificates are present and valid by accessing the controller's command-line interface and issuing the command **show tpm cert-info**. If the controller has a valid certificate, the output of the command appears similar to the output in the example below.

```
(host) # show tpm cert-info
subject= /CN=AC1234567::00:0b:86:11:22:33
issuer= /DC=com/DC=companyname/DC=ca3/CN=DEVICE-CA3
serial=5147D5BC000000000000C
notBefore=Aug 29 22:16:12 2009 GMT
notAfter=Aug 18 22:16:12 2029 GMT
```

If the controller displays the following output, it may have a corrupted or missing TPM and factory certificates. Contact Dell support.

```
(host) # show tpm cert-info
Cannot get TPM and Factory Certificate Info.
```

Disabling Control Plane Security

If you disable control plane security on a standalone or local controller, all APs connected to that controller reboot then reconnect to the controller over a clear channel.

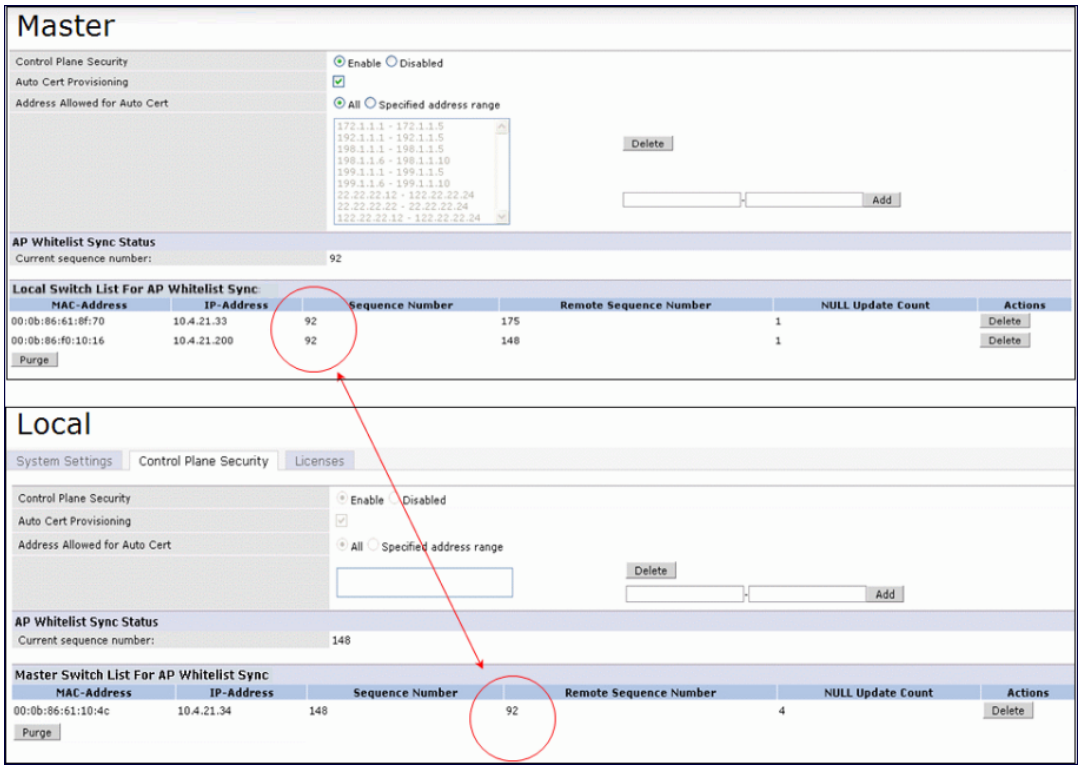
If your disable control plane security on a *master* controller, APs directly connected to the master controller reboot then reconnect to the master controller over a clear channel. However, its local controllers continue to communicate with their APs over a secure channel until you save your configuration on the master controller. Once you save the configuration, the changes are pushed down to the local controllers. At that point, any APs connected to the local controllers also reboot and reconnect over a secure channel.

Verifying Whitelist Synchronization

To verify that a network of master and local controllers are correctly sharing their campus AP whitelists, check the sequence numbers on the master and local controller whitelists.

- The sequence number value on a master controller should be the same as the remote sequence number on the local controller.
- The sequence number value on a local controller should be the same as the remote sequence number on the master controller.

Figure 8 *Sequence numbers on Master and Local Controllers*



Rogue APs

If you enable auto certificate provisioning enabled with the **Auto Cert Allow All** option, any AP that appears on the network receives a certificate. If you notice unwanted or rogue APs connecting to your controller via an IPsec tunnel, verify that automatic certificate provisioning has been disabled, then manually remove the unwanted APs by deleting their entries from the campus AP whitelist.

ArubaOS base features include sophisticated authentication and encryption, protection against rogue wireless APs, seamless mobility with fast roaming, the origination and termination of IPsec/L2TP/PPTP tunnels between controllers, clients, and other VPN gateways, adaptive RF management and analysis tools, centralized configuration, and location tracking.

Optional add-on licenses provide advanced feature such as Wireless Intrusion Protection and Policy Enforcement Firewall. Evaluation licenses are available for some of these advanced features.

ArubaOS licenses are detailed in the following sections:

- Understanding License Terminology on page 130
- Working with Licenses on page 131
- Centralized Licensing in a Multi-Controller Network on page 132
- Using Licenses on page 142
- License Installation Best Practices and Exceptions on page 144
- Installing a License on page 144
- Deleting a License on page 146
- Moving Licenses on page 147
- Resetting the Controller on page 147

Understanding License Terminology

For clarity, the following terminology is used throughout this chapter.

- Bundle: a cost-effective way to purchase functionality that supports a controller and *x*-number of APs.
- Certificate ID: the identification number attached to the Software License Certificate. The Certificate ID is used in conjunction with the controller's serial number to create the License Key.
- Evaluation License: a license that allows you to evaluate a feature set (or module) for a maximum of 90 days. The evaluation licenses are uploaded in 30-day increments. Only modules that offer new and unique functionality support Evaluation Licenses.
- License Certificate: a certificate (soft copy) that contains license information including:
  - License Description
  - Quantity
  - Part Number/Order Number
  - Certificate ID
- License Database: the licenses installed on your controller
- License Key: generated from the controller serial number
- Permanent License: the opposite of an evaluation license. This license permanently installs the specific features represented by the license.
- Upgrade License: a license that adds AP capacity to your controller. Note that Upgrade Licenses do not support an evaluation license.

Working with Licenses

Each license refers to specific functionality (or module) that supports unique features.

The licenses are:

- Base OS: base operating functions including VPN and VIA clients.
- AP Capacity: capacity license for RAP indoor and outdoor Mesh APs. Campus, Remote, or Mesh APs can terminate on the controller without the need for a separate license.
- Advanced Cryptography (ACR): this is required for the Suite B Cryptography in IPsec and 802.11 modes. License enforcement behavior controls the total number of concurrent connections (IPsec or 802.11) using Suite B Cryptography.The xSec license features are bundled with this license.
- Policy Enforcement Firewall Virtual Private Network (PEFV): enables Policy Enforcement Firewall for VIA clients. This is a controller license.
- Policy Enforcement Firewall Next Generation (PEFNG): Wired, WLAN Licensed per AP numbers including user roles, access rights, Layers 4 through 7 traffic control, per-service prioritization/QoS, authentication/accounting APIs, External Service Interfaces (ESI), Voice and Video. This is an AP count license.
- Public Access: reserved for future use.
- RFProtect: Wireless Intrusion Protection (WIPS) and Spectrum Analysis. This is an AP count license.
- xSec (Extreme Security) for Federal: Layer 2 VPN for wired or wireless using FIPS-approved algorithms.
- Internal Test Functions: for internal use only.

The license categories are:

- Permanent license: this type of license permanently enables the desired software module on a specific Dell controller. You obtain permanent licenses through the sales order process only. Permanent software license keys are sent to you via email.
- Evaluation license: this type of license allows you to evaluate the unrestricted functionality of a software module on a specific controller for 90 days (in three 30-day increments).

  An expired evaluation license will remain in the license database until the controller is reset using the command **write erase all** where all license keys are removed. An expired evaluation license has no impact on the normal operation of the controller. It is kept in the license database to prevent abuse.



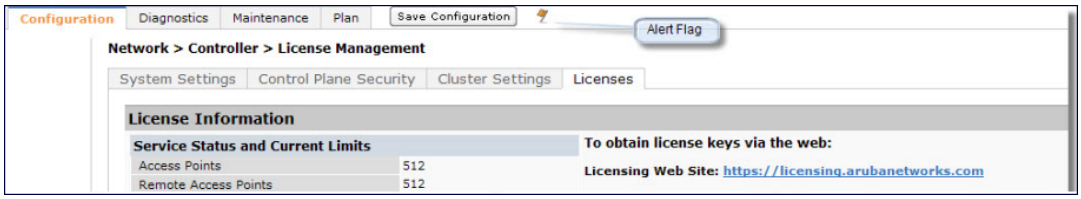
When you apply license keys on a controller, abnormal tampering of the device's system clock (setting the system clock back) results in the disabling of software licensed modules and their supported features. This can affect network services.

To determine your time remaining on an evaluation license, a banner is displayed when you log in through the command line:

```
NOTICE
NOTICE -- This switch has active licenses that will expire in 29 days
NOTICE
NOTICE -- See 'show license' for details.
NOTICE
```

From the WebUI, an "Alert" appears with information regarding the evaluation license status (see Figure 9).

Figure 9 *Alert Flag*



At the end of the 90-day period, you must apply for a permanent license to re-enable the features permanently on the controller. Evaluation software license keys are only available in electronic form and are emailed to you.

When an evaluation period expires:

- The controller automatically backs up the startup configuration and reboots itself at midnight (according to the system clock).
- All permanent licenses are unaffected. The expired evaluation licensed feature is no longer available and is displayed as **Expired** in the WebUI.

- **Upgrade license**—This license expands AP capacity. There are no Evaluation licenses available for Upgrade licenses.

Centralized Licensing in a Multi-Controller Network

In order to configure each feature on the local controller, the master controller(s) must be licensed for each feature configured on the local controllers. Centralized licensing simplifies licensing management by distributing licenses installed on one controller to other controllers on the network. One controller acts as a centralized license database for all other controllers connected to it, allowing all controllers to share a pool of unused licenses. The primary and backup licensing servers can share a single set of licenses, eliminating the need for a redundant license set on the backup server. Local licensing client controllers maintain information sent from the licensing server, even if the licensing client controller and the licensing server controller can no longer communicate. If an AP fails over from one client controller to another, the AP will be allowed to come up even if there aren't sufficient licenses present on the backup controller. the APs continue to stay active until they reboot. However, if there are not sufficient available licenses to bring up an AP after it reboots, that AP will not become active.

You can use the centralized licensing feature in a master-local topology with a redundant backup master, or in a multi-master network where all the masters can communicate with each other (for example, if they are all connected to a single AirWave server). In the master-local topology, the master controller acts as the primary licensing server, and the redundant backup master acts as the backup licensing server. In a multi-master network, one controller must be designated as a primary server, and a second controller must be configured as a backup licensing server.

Centralized licensing can distribute the following license types:

- AP
- PEFNG
- RFProtect
- xSec
- ACR

This section includes the following topics:

- Primary and Backup Licensing Servers
- Communication between the License Server and License Clients

- Replacing a Controller
- Failover Behaviors
- Configuring Centralized Licensing

Primary and Backup Licensing Servers

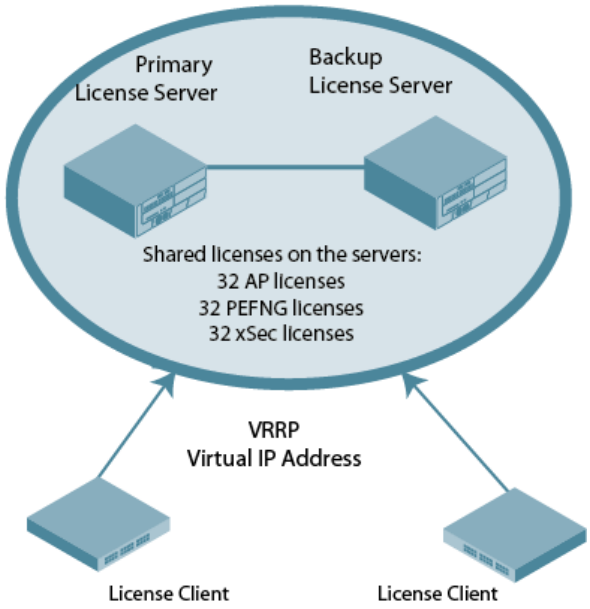
Centralized licensing allows the primary and backup licensing server controllers to share a single set of licenses. If you do not enable this feature, the master and backup master controller each require separate, identical license sets. The two controllers acting as primary and backup license servers must use the same version of ArubaOS, and must be connected on the same broadcast domain using the Virtual Router Redundancy Protocol (VRRP). Other client controllers on the network connect to the licensing server using the VRRP virtual IP address configured for that set of redundant servers. The primary licensing server uses the configured virtual IP address by default. However, if the controller acting as the primary licensing server becomes unavailable, the secondary licensing server will take ownership of the virtual IP address, allowing licensing clients to retain seamless connectivity to a licensing server.



Only one backup licensing server can be defined for each primary server.

The example below shows a primary and backup license server connected using VRRP. Licenses installed on either the primary or the backup server are shared between that pair of servers. If the primary and backup controllers each had 16 AP licenses, 16 PEFNG licenses, and 16 xSec licenses installed, they would share a combined pool of 32 AP, 32 PEFNG, and 32 xSec licenses. Any license client controllers connected to this pair of redundant servers could also use licenses from this license pool.

Figure 10 *Shared Licenses on a Primary and Backup Licensing Server*



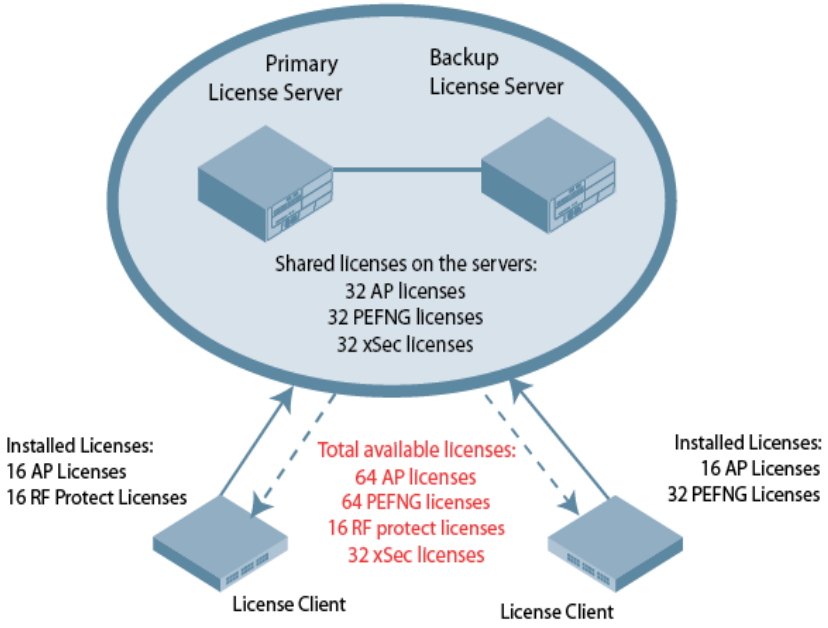
Communication between the License Server and License Clients

When you enable centralized licensing, information about the licenses already installed on the individual client controllers are sent to the licensing server, where they are added into the server's licensing table. The information in this table is then shared with all client controllers as a pool of available licenses. When a client controller uses a license in the available pool, it communicates this change to the licensing server master controller, which updates the table before synchronizing it with the other clients.

Client controllers do not share information about built-in licenses to the licensing server. A controller using the centralized licensing feature will use its built-in licenses before it consumes available licenses from the license pool. As a result, when a client controller sends the licensing server information about the licenses that a client is using, it only reports licenses taken from the licensing pool, and disregards any built-in licenses used. For example, if a controller has a built-in 16-AP license and twenty connected APs, it will disregard the built-in licenses being used, and will report to the licensing server that it is using only four AP licenses from the license pool.

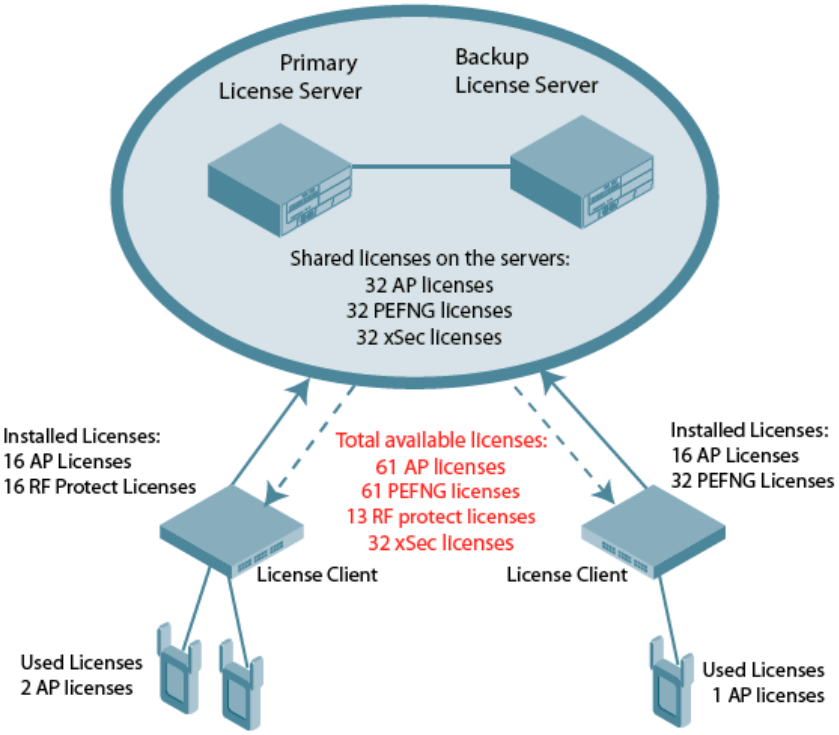
When centralized licensing is first enabled on the licensing server, its licensing table only contains information about the licenses installed on that server. When the clients contact the server, the licensing server adds the client licenses to the licensing table, then sends the clients information about the total available licenses for each license type. In the following example, the licenses installed on two client controllers are imported into the license table on the license server. The licensing server then shares the total number of available licenses with other controllers on the network.

Figure 11 *Licenses Shared by Licensing Clients*



When new AP associates with a licensing client, the client sends updated licensing information to the server. The licensing server then recalculates the available total, and sends the revised license count back to the clients. If a client uses an AP license from the license pool, it also consumes a PEFNG and a RFProtect license from the pool, even if that AP has not enabled any features that would require that license. A controller cannot use more licenses than what is supported by its controller platform, regardless of how many licenses are available in the license pool.

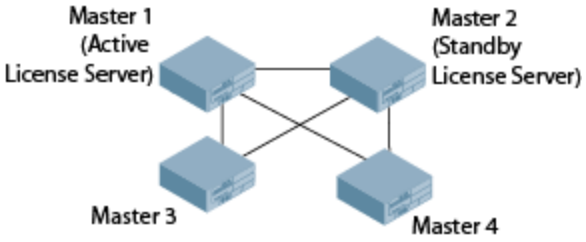
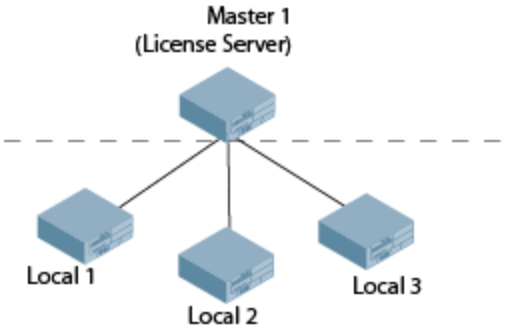
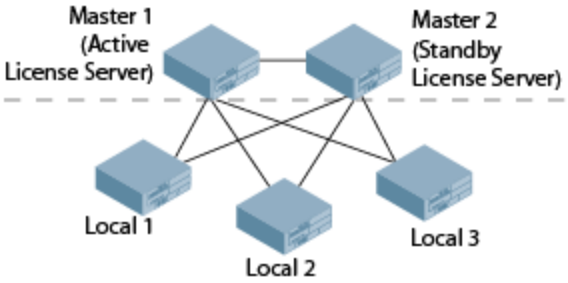
Figure 12 *License Pool Reflecting Used licenses*



Supported Topologies

The following table describes the controller topologies supported by this feature.

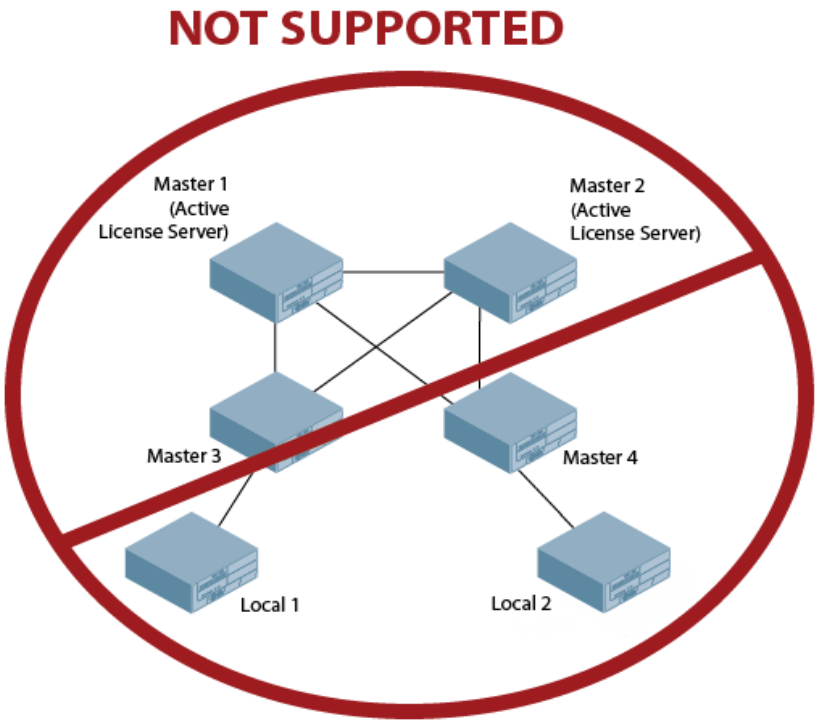
Table 21: *Centralized Licensing Topologies*

Topology	Example
All controllers are master controllers. The master and standby licensing servers must be defined.	
A single master controller is connected to one or more local controllers. Only the master controller can be a license server. A local controller can only be license client, not a license server.	
A master and standby master are connected to one or more local controllers. The master license server will reside on the master controller, and the standby license server will reside on the standby master controller. Local controllers can only be license clients, not license servers.	

Unsupported Topologies

The centralized licensing feature does NOT support topologies where multiple master controllers have one or more attached local controllers.

Figure 13 *Topologies Not Supported by Centralized Licensing*



Adding and Deleting Licenses

New licenses can be added to any controller managed by a centralized licensing system, although best practices recommend adding them to the primary licensing server for easier management and tracking of licenses across a wide network. Licenses can only be deleted from the controller on which the license is installed.

You do not need to reboot a controller after adding or deleting a license, regardless of whether you enable centralized licensing. If you delete a license from a licensing client or server and there are no longer enough licenses to support the number of active APs on the network, the APs continue to stay active until they reboot. If there are not sufficient available licenses to bring up an AP after it reboots, that AP will not become active.

Centralized licensing supports evaluation licenses. When a client controller has an evaluation license installed, those license limits will be sent to the licensing server and added to the license pool as long as the evaluation period is active. When the evaluation period expires, the client with the expired license sends its revised limits to the license server. The licensing server removes the evaluation licenses from its license table, then sends updated license pool information to other clients on the network.

Replacing a Controller

If you need to replace the controller acting as a license server, the keys installed on the previous license server must be regenerated and added to the new license server. If you need to replace a controller acting as license client, you must regenerate the license keys installed on the client and reinstall them on the replacement client or the licensing server.

Failover Behaviors

If the primary licensing server fails, the controller acting as a backup license server will retain the shared license limits until the backup server reboots. If both the primary and the backup license servers fail, or if the backup controller reboots before the primary controller comes back up, License clients will retain the license limits sent to them by the licensing server for 30 days.



Although a client controller retains its licensing information for 30 days after it loses contact with the licensing server, if the client reboots at any time during this 30-day window, the window will restart, and the client will retain its information for another 30 days.

APs that use centralized licensing in conjunction with a ArubaOS high availability feature behave differently than APs that do not use a high availability solution. APs using VRRP redundancy, a backup LMS, or the ArubaOS fast failover feature can quickly fail over to a backup controller, even if that backup controller does not have any AP licenses at the time of the failover. However, if that AP reboots, it will not obtain its licenses until the backup controller receives the required licenses from the licensing master.

Client is Unreachable

The centralized licensing feature sends keepalive heartbeats between the license server and the licensing client controllers every 30 seconds. If the licensing server fails to receive three consecutive heartbeats from a client, it assumes that the licensing client is down, and that any APs associated with that client are also down or have failed over to another controller . Therefore, the licensing server adds any licenses used by that client back into to the available pool of licenses. If the license server fails to contact a license client for 30 consecutive days, any licenses individually installed on that client will be removed from the server's license database.



The WebUI of the licensing client and the licensing server both display a warning message when a licensing client and licensing server are unable to communicate.

Server is Unreachable

If a licensing client does not receive three consecutive heartbeats from the server, it assumes that the server is down, and that any APs directly associated to the server are also down or have failed over to another controller. The client then adds any licenses used by the licensing server into to the pool of available licenses on that client. When a license client is unable to reach a license server for 30 consecutive days, it removes any shared licenses pushed to it from the licensing server, and reverts to its installed licenses. If the 30-day window has passed and the controller does not have enough installed licenses for all of its associated APs, the controller will nonetheless continue to support each AP. However, when an AP reboots and its controller does not have enough licenses, that AP will not come up.

Configuring Centralized Licensing

The steps to configure centralized licensing on your network vary, depending upon whether you are enabling this feature in a network with a master-local controller topology, or in a network where all controllers are configured as masters. Before you enable this feature, you must ensure that the controllers are able to properly communicate with the licensing master. Once you have identified your deployment type, follow the steps in the appropriate section below.

Pre-configuration Setup in an All-Master Deployment

Follow the steps described below to configure the centralized licensing feature in a network with all master controllers.

1. Ensure that the controllers that will use this feature are associated with the same AirWave server.
2. Identify a controller you want to designate as the primary licensing server. If that controller already has a redundant backup controller, that backup controller will automatically become the backup license server.
3. (Optional) If your primary licensing server does not yet have a dedicated, redundant backup controller and you want to use a backup server with the centralized licensing feature, you must identify a second controller to use as the backup licensing server, and create a virtual router on the primary licensing server.
4. (Optional) Establish secure IPsec tunnels between the primary licensing server controller and the licensing client controllers by enabling control plane security on that cluster of master controllers, or by creating site-

to-site VPN tunnels between the licensing server and client controllers. This step is not required, but if you do not create secure tunnels between the controllers, the controllers will exchange clear, unencrypted licensing information. This step is not required for a master-local topology.

Preconfiguration Setup in a Master/Local Topology

The master controller in a master-local topology is the primary licensing server by default. If this master controller already has a redundant standby master, that redundant master will automatically act as the backup licensing server with no additional configuration. If your primary licensing server does not yet have a redundant standby controller and you want to use a backup server with the centralized licensing feature, you must identify a second controller you want to designate as the backup licensing server, and define a virtual router on the primary licensing server.

Enabling Centralized Licensing

The following steps describe the procedure to enable centralized licensing on both the licensing master and the licensing clients.

Using the WebUI

1. Access the WebUI of the primary licensing master controller, navigate to **Configuration > Controller** and select the **Centralized Licenses** tab.
2. Select **Enable Centralized Licensing**.
3. (Optional) If the licensing server already has a dedicated redundant standby controller, that standby controller will automatically become the backup license server. If the primary licensing server in your deployment does not have a dedicated, redundant master controller, but you want to define a backup server for the licensing feature, follow steps a-c below:
   a. In the **VRRP ID** field, enter the Virtual Router ID for the Virtual Router you configured in the Preconfiguration Setup task in the section above.
   b. In the **Peer's IP address** field, enter the IP address of the backup licensing server.
   c. In the **License Server IP** field, enter the virtual IP address for the Virtual Router used for license server redundancy.
4. Click **Apply**.

If you are deploying centralized licensing on a cluster of master controllers, you must define the IP address that the licensing clients in the cluster use to access the licensing server.

5. Access the WebUI of a licensing client, navigate to **Configuration > Controller** and select the **Centralized Licenses** tab.
6. Select **Enable Centralized Licensing**.
7. In the License Server IP field, enter the IP address the client will use to connect to the licensing server. If you have defined a backup licensing server using a virtual router ID, enter the IP address of that virtual router.
8. Click **Apply**.
9. Repeat steps 5-8 on each licensing client in the cluster.

Using the CLI

Access the command-line interface of the licensing server, and issue the following commands in config mode:

```
(host)(config) #license profile
(host)(License provisioning profile) #centralized-licensing-enable
```

If the licensing server already has a dedicated redundant standby controller, that standby controller will automatically become the backup license server. If the primary licensing server in your deployment does not

have a redundant master controller but you want to define a backup server for the licensing feature, issue the following commands on the licensing server:

```
(host)(License provisioning profile) #License server-redundancy
(host)(License provisioning profile) #License-vrrp <vrId>
(host)(License provisioning profile) #Peer-ip-address <ip>
```

If you are deploying centralized licensing on a cluster of master controllers, access the command-line interface of a licensing client controller, and issue the following commands in config mode:

```
(host) (config) #license profile
(host) (License provisioning profile) #centralized-licensing-enable
(host) (License provisioning profile) #license server-ip <ip>
```

If a controller is designated as standby license server, it does not have the **license-server-ip** value configured.

Monitoring and Managing Centralized Licenses

A centralized licensing server displays a wide variety of licensing data that you can use to monitor licenses and license usage. The tables described below are available on the **Network > Controller > Centralized License Management > Information** page of the Licensing server WebUI.

License server Table

This table displays information about the different types of licenses in the license table, and how many total licenses of each type are available and used. This table includes the following information:

Table 22: *License Server Table Data*

Column	Description
Service Type	Type of license on the licensing server.
Aggregate Licenses	Number of licenses in the licensing table on the licensing server.
Used Licenses	Total number of licenses of each license type reported as used by the licensing clients or licensing server.
Remaining Licenses	Total number of remaining licenses available in the licensing table.

License Client Table

This table displays centralized license limits applied to each licensing client. This table includes the following information:

Table 23: *License Client Table Data*

Column	Description
Service Type	Type of license on the licensing client.
System Limit	The maximum number of licenses supported by the controller platform.
Server Licenses	Number of licenses sent from the licensing server.. **NOTE:** This number is limited by the total license capacity of the controller platform. A controller cannot use more licenses than is supported by that controller platform, even if additional license are available.

Column	Description
Used Licenses	Total number of licenses of each license type used by the licensing client controller.
Contributed Licenses	Total number of licenses of each license type contributed by the licensing client controller.
Remaining Licenses	Total number of remaining licensing available on this controller. This number is also limited by the total license capacity of the controller platform.

License Client(s) Usage Table

This table displays information about the different types of licenses in the license table, and how many total licenses of each type are available and used.

Table 24: *License Clients(s) Usage Table Data*

Column	Description
Hostname	Name of the licensing client controller.
IP Address	IP address of the licensing client controller.
AP	Total number of AP licenses used by a licensing client associated with this controller.
PEF	Total number of Policy Enforcement Firewall (PEF) licenses used by a licensing client associated with this controller.
RF Protect	Total number of RFProtect licenses used by a licensing client associated with this controller.
xSec Module	Total number of Extreme Security (xSec) licenses used by a licensing client associated with this controller.
ACR	Total number of advanced Cryptography (ACR) licenses used by a licensing client associated with this controller.
Last update (secs. ago)	Time, in seconds, that has elapsed since the licensing client received a heart-beat response.

Aggregate License Table

Issue this command from the command-line interface of the centralized licensing server controller to view license limits sent by licensing clients.

Table 25: *Aggregate License Table Data*

Column	Description
Hostname	Name of the licensing client controller.
IP Address	IP address of the licensing client controller.
AP	Total number of AP licenses sent from licensing clients associated with this controller.
PEF	Total number of Policy Enforcement Firewall (PEF) licenses sent from licensing clients associated with this controller.
RF Protect	Total number of RFProtect licenses sent from licensing clients associated with this controller.
xSec Module	Total number of Extreme Security (xSec) licenses sent from licensing clients associated with this controller.
ACR	Total number of advanced Cryptography (ACR) licenses sent from licensing clients associated with this controller.

License Heartbeat Table

This table displays the license heartbeat statistics between the license server and the license client.

Table 26: *License Heartbeat Table Data*

Column	Description
IP address	IP address of the licensing client.
HB Req	Heartbeat requests sent from the licensing client.
HB Resp	Heartbeat responses received from the license server.
Total Missed	Total number of heartbeats that were not received by the licensing client.
Last Update	Number of seconds elapsed since the licensing client last sent a heartbeat request.

Using Licenses

Licenses are platform independent and can be installed on any controller. Installation of the feature license unlocks that feature's functionality for the maximum capacity of the controller.



The license limits are enforced until you reach the controller limit (see Table 28).

Table 27 lists how licenses are consumed on the Controllers.

Table 27: *Usage per License*

License	Basis	What Consumes One License
PEFNG	AP	One operational AP
xSec	Session	One active client termination
RFprotect	AP	One operational AP
AP	AP	One operational LAN-connected or mesh AP that is advertising at least one BSSID (virtual-AP) or RAP
ACR	Session	One active client termination

The controller licenses are variable-capacity (see Table 28).



In Table 28, the Remote AP count is equal to the total AP count for all the controllers. The Campus AP count is 1/4 of the total AP count *except* for the W-6000M3 which is one half the AP count.

Table 28: *Controller AP Capacity*

Controller	Total AP Count	Campus APs	Remote APs
W-7210	512	512	512
W-7220	1024	1024	1024
W-7240	2048	2048	2048
W-6000M3	1024	512	1024
W-3200	128	32	128
W-3400	256	64	256
W-3600	512	128	512
W-620	8	8	8
W-650	16	16	16

Understanding License Interaction

The some licenses interact with each other, and may require some equality.

- AP/PEFNG and RFProtect must be equal.
  - All active APs run AP/PEFNG and RFProtect services (if enabled). If they are not equal, the number of active APs are restricted to the minimum of the AP/PEFNG and RFProtect license count.



It is not possible to designate specific APs for RFProtect/non-RFProtect operations.

  - Mesh portals/mesh points with no virtual APs, do not consume am RFProtect license

- If a Mesh node is also configured for client service (for example, it advertises a BSSID), it consumes one AP license.
- Remote APs consume licenses the same as campus APs.
- ACR Interaction
  - On a platform that supports 2048 IPsec tunnels, the maximum number of Suite B IPsec tunnels supported is 2048, even if a larger capacity license is installed.
  - The ACR license is cumulative. If you want to support 2048 Suite B connections, install two ACR licenses (LIC-ACR-1024).
  - An evaluation ACR license is available (EVL-ACR-1024). You can install the ACR evaluation license with a higher capacity than the platform maximum.
  - On a platform that supports 2048 IPsec tunnels, with a LIC-ACR-512 installed, only 512 IPsec tunnels can be terminated using Suite B encryption. An additional 1536 IPsec tunnels, using non-Suite B modes (for example, AES-CBC), can still be supported.
  - On a platform with LIC-ACR-512 installed, a mixture of IPsec and 802.11i Suite B connections can be supported. The combined number of these sessions may not exceed 512.
  - A single client using both 802.11i Suite B and IPsec Suite B simultaneously will consume two ACR licenses.

License Installation Best Practices and Exceptions

- Back up the controller's configuration (**backup flash** command) and back up the License database (**license export *filename***) before making any changes.

```
(host) #backup flash
Please wait while we tar relevant files from flash...
Please wait while we compress the tar file...
Checking for free space on flash...
Copying file to flash...
File flashbackup.tar.gz created successfully on flash.
Please copy it out of the switch and delete it when done.
(host) #license export licensebackup.db
Successfully exported 1 licenses from the License Database to licensebackup.db
```

- Allow for the maximum quantity required at any given time.
- When calculating AP licenses, determine the normal AP load of your controller and add a backup load for failure scenarios.
- Use 20 users per AP as a reasonable estimate when calculating user licenses. Do not forget to consider occasional large assemblies or gatherings.

Installing a License

The Dell licensing system is controller-based. A license key is a unique alphanumerical string generated using the controller's serial number and is valid only for that controller only. Licenses can be pre-installed at the factory so that all licensed features are available upon initial setup. You can also install license features yourself.



It is recommended that you obtain a user account on the Dell Software License Management website even if software license keys are preinstalled on your controller.

Enabling a new license on your controller

The basic steps to installing and enabling a new license feature are listed below along with a reference to a section in this document with more detailed information.

1. Obtain a valid Dell software license from your sales account manager or authorized reseller (see Requesting a Software License in Email on page 145).
2. Locate the system serial number of your controller (see Locating the System Serial Number on page 145).
3. Use your system's serial number to obtain a software license key from the Dell Software License Management website.licensing.dell-pcw.com (see Obtaining a Software License Key on page 145).
4. Enter the software license key via the controller's WebUI using one of the following procedures
   - navigate to **Configuration > Network > Controller > System Settings** page and select the **License** tab. Enter the software license key and click **Apply** (see Applying the Software License Key in the WebUI on page 146).
   - Launch the License Wizard from the **Configuration** tab and click **New**. Enter the software license key in the space provided (see Applying the Software License Key in the License Wizard on page 146).

Requesting a Software License in Email

To obtain either a permanent or a evaluation software license, contact your sales account manager or authorized reseller. The license details are provided via email with an attached text file. Use the text file to cut and paste the licensing information into the WebUI or at the command line.



Ensure that you have provided your sales person with a valid email address.

The email also includes:

- The orderable part number for the license
- A description of the software module type and controller for which it is valid
- A unique, 32-character alphanumerical string used to access the license management website and which, in conjunction with the serial number of your controller, generates a unique software license key

Locating the System Serial Number

Each controller has a unique serial number located at the rear of the controller chassis. The W-6000M3 has the serial number on the device itself.

You can also find the serial numbers by navigating to the **Controller > Inventory** page on the WebUI or by executing the **show inventory** command from the CLI.



To physically inspect the system serial number on a W-6000M3 , you need to remove the device from the controller chassis, which may result in network down time.

Obtaining a Software License Key

To obtain a software license key, you must log in to the Dell License Management website. If you are a first time user of the licensing site, you can use the software license certificate ID number to log in initially and request a user account. If you already have a user account, log in to the site.

Once logged in, you are presented with several options:

- **Activate a certificate**: Activate a new certificate and create the software license key that you will apply to your controller.

- **Transfer a certificate**: Transfer a software license certificate ID from one controller to another (for example, transferring licenses to a spare system).
- **Import preloaded certificates**: For controllers on which licenses are pre-installed at the factory. transfer all software license certificate IDs used on the sales order to this user account.
- **List your certificates**: View all currently available and active software license certificates for your account.

Creating a Software License Key

To create a software license key, you must log to to the Dell License Management website at:

licensing.dell-pcw.com/

If you are a first time user of the licensing site, you can use the software license certificate ID number to log in initially and request a user account. If you already have a user account, log in to the site.

1. Select **Activate a Certificate.**
2. Enter the certificate ID number and the system serial number of your controller.
3. Review the license agreement and select **Yes** to accept the agreement.
4. Click **Activate it**. A copy of the transaction and the software license key is emailed to you at the email address you entered for your user account



The software license key is valid *only* for the system serial number for which you activated the certificate.

Applying the Software License Key in the WebUI

To enable the software module and functionality, you must apply the software license key to your controller.

1. Log in to your controller's WebUI.
2. Navigate to the **Configuration > Network > Controller > System Settings** page and select the **License** tab.
3. Copy the software license key, from your email, and paste it into the **Add New License Key** field.
4. Click **Add**.

Applying the Software License Key in the License Wizard

Log in to your controller's WebUI.

1. Launch the License Wizard from the **Configuration** tab and click **New** .
2. The License Wizard help walk you through the activation process. Click the **Help** tab within the License Wizard for additional assistance.

Deleting a License

To remove a license from a system:

1. Navigate to the **Configuration > Network > Controller > System Settings** page and select the **License** tab.
2. Scroll down to the **License Table** and locate the license you want to delete.
3. Click **Delete** at the far right hand side of the license to delete the license.

   If a license feature is under an evaluation license, it will not generate a key is generated when the feature is deleted.

Moving Licenses

It may be necessary to move licenses from one controller to another or to delete a license for future use. To move licenses, delete the license from the chassis as described in Deleting a License on page 146. Then install the license key on the new controller as described in Applying the Software License Key in the WebUI on page 146.



ArubaOS provides the ability to move a license from one controller to another, for maximum flexibility in managing an organization's network and to minimize an RMA impact. Dell monitors and detects license fraud. Abnormally high volumes of license transfers for the same license certificate to multiple controllers can indicate a breach of the Dell end user software license agreement and will be investigated.

Resetting the Controller

Rebooting or resetting a controller has no effect on either a permanent or a evaluation license.

Issuing the **write erase** command on a controller running software licenses does *not* affect the license key management database on the controller.

Issuing the **write erase all** command resets the controller to factory defaults, and deletes all databases on the controller, including the license key management database. You must reinstall all previously-installed license keys.

On a W-7200 Series controller, you can reset controller using the LCD screen. Issuing the **Factory Default** option under the **Maintenance** menu returns the controller to the factory default settings. For more information about the LCD menu, see Using the LCD Screen on page 99.

The following topics in this chapter describe some basic network configuration on the controller:

- Configuring VLANs on page 148
- Configuring Ports on page 155
- Understanding VLAN Assignments on page 157
- Configuring Static Routes on page 165
- Configuring the Loopback IP Address on page 165
- Configuring the Controller IP Address on page 166
- Configuring GRE Tunnels on page 167
- Jumbo Frame Support on page 171

Configuring VLANs

The controller operates as a layer-2 switch that uses a VLAN as a broadcast domain. As a layer-2 switch, the controller requires an external router to route traffic between VLANs. The controller can also operate as a layer-3 switch that can route traffic between VLANs defined on the controller.

You can configure one or more physical ports on the controller to be members of a VLAN. Additionally, each wireless client association constitutes a connection to a *virtual port on the controller*, with membership in a specified VLAN. You can place all authenticated wireless users into a single VLAN or into different VLANs, depending upon your network. VLANs can remain inside the controller, or they can extend outside the controller through 802.1q VLAN tagging.

You can optionally configure an IP address and netmask for a VLAN on the controller. The IP address is *up* when at least one physical port in the VLAN is up. The VLAN IP address can be used as a gateway by external devices; packets directed to a VLAN IP address that are not destined for the controller are forwarded according to the controller's IP routing table.

Creating and Updating VLANs

You can create and update a single VLAN or bulk VLANs.

In the WebUI

1. Navigate to the **Configuration > Network > VLANs** page.
2. Click **Add a VLAN** to create a new VLAN. (To edit an existing VLAN, click **Edit** for the VLAN entry.) See Creating Bulk VLANs In the WebUI on page 149 to create a range of VLANs.
3. In the **VLAN ID** field, enter a valid VLAN ID. (Valid values are from 1 to 4094, inclusive).
4. To add physical ports to the VLAN, select **Port**. To associate the VLAN with specific port-channels, select **Port-Channel**.
5. (Optional) Click the **Wired AAA Profile** drop-down list to assign an AAA profile to a VLAN. This wired AAA profile enables role-based access for wired clients connected to an untrusted VLAN or port on the controller.

   Note that this profile will only take effect if the VLAN or port on the controller is untrusted. If you do not assign a wired AAA profile to the VLAN, the global wired AAA profile applies to traffic from untrusted wired ports.

6. If you selected **Port** in step 4, select the ports you want to associate with the VLAN from the **Port Selection** window.

   or

   If you selected **Port-Channel** in step 4, click the **Port-Channel ID** drop-down list, select the specific channel number you want to associate with the VLAN, then select the ports from the **Port Selection** window.
7. Click **Apply**.

In the CLI

Use the following commands:

```
(host)(config) #vlan <id>
(host)(config) #interface fastethernet|gigabitethernet <slot>/<port>
(host)(config-if) #switchport access vlan <id>
```

Creating Bulk VLANs In the WebUI

1. To add multiple VLANs at one time, click **Add Bulk VLANs**.
2. In the **VLAN Range** pop-up window, enter a range of VLANs you want to create at once. For example, to add VLAN IDs numbered 200-300 and 302-350, enter 200-300, 302-350.
3. Click **OK**.
4. To add physical ports to a VLAN, click **Edit** next to the VLAN you want to configure and click the port in the **Port Selection** section.
5. Click **Apply**.

In the CLI

Use the following commands:

```
(host)(config) #vlan
(host)(config) #vlan range 200-300,302-350
```

Creating a VLAN Pool

You can create, update, and delete a VLAN pool. Each VLAN pool has a name and needs to have one or more VLANs assigned to it. The following configurations create a VLAN Pool named **mygroup**. It has the assignment type **Even**, and VLAN IDs 2, 4 and 12 are assigned to this pool.



ArubaOS supports maximum of 256 VLANs per VLAN Pool.

Using the WebUI

1. Navigate to **Configuration > Network > VLANs.**
2. Select the **VLAN Pool** tab to open the **VLAN Pool** window.
3. Click **Add**.
4. In the **VLAN Name** field, enter a name that identifies this VLAN pool.
5. In the **Assignment Type** field, select **Hash** or **Even** from the drop-down list. See Distinguishing Between Even and Hash Assignment Types on page 150 for information and conditions regarding Hash and Even assignment types.

The Even VLAN pool assignment type is only supported in tunnel and dtunnel modes. It is not supported in split or bridge modes. It is not allowed for VLAN pools that are configured directly under a virtual AP (VAP). It must only be used under named VLANs. L2 Mobility is not compatible with the existing implementation of the Even VLAN pool assignment type.

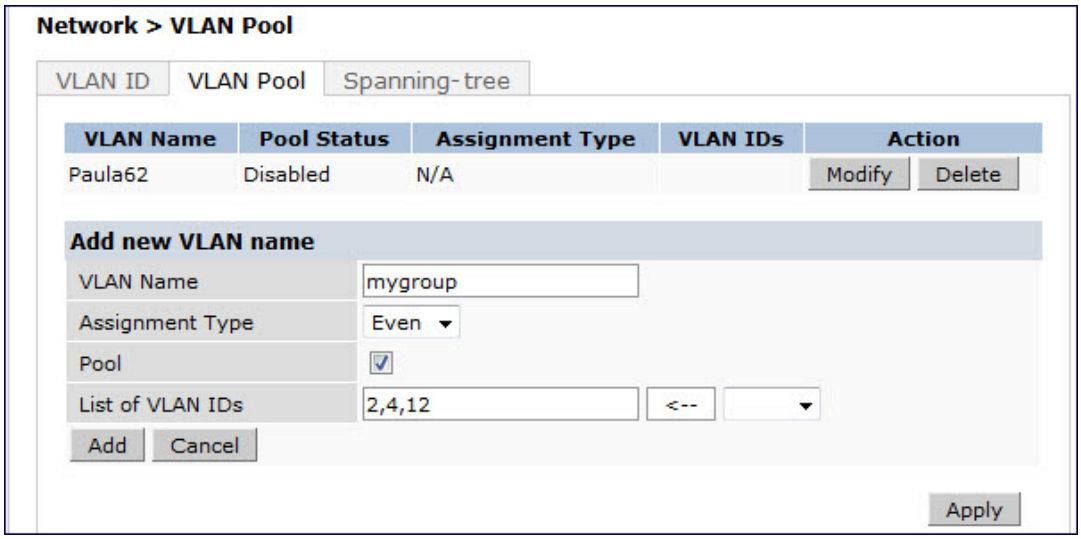
6. Check the **Pool** check box if you want the VLAN to be part of a pool.
7. In the **List of VLAN IDs** field, enter the VLAN IDs you want to add to this pool. If you know the ID, enter each ID separated by a comma. You can also click the drop-down list to view the IDs, then click the <-- arrow to add the ID to the pool.

VLAN pooling should *not* be used with static IP addresses.

8. You must add two or more VLAN IDs to create a pool.
9. When you finish adding all the IDs, click **Add**.

   The VLAN pool along with its assigned IDs appears on the VLAN Pool window. If the pool is valid, its status is enabled.

Figure 14 *Creating a VLAN Pool*



10. Click **Apply**.
11. At the top of the window, click **Save Configuration**.

Distinguishing Between Even and Hash Assignment Types

The VLAN assignment type determines how the controller handles a VLAN assignment.

The Hash assignment type means that the VLAN assignment is based on the station MAC address. The Even assignment type is based on an even distribution of VLAN pool assignments.

The Even VLAN Pool assignment type maintains a dynamic latest usage level of each VLAN ID in the pool. Therefore, as users age out, the number of available addresses increases. This leads to a more even distribution of addresses.

The Even type is only supported in tunnel and dtunnel modes. It is not supported in split or bridge modes and it is not allowed for VLAN pools that are configured directly under a virtual AP. It can only be used under named VLANs.

If a VLAN pool is given an Even assignment and is assigned to user roles, user rules, VSA, or server derivation rules, then while applying VLAN derivation for the client "on run time," the Even assignment is ignored and the Hash assignment is applied with a message displaying this change.



L2 Mobility is not compatible with the existing implementation of the Even VLAN pool assignment type.

Updating a VLAN Pool

1. On the **VLAN Pool** window, click **Modify** next to the VLAN name you want to edit.
2. Modify the assighment type and the list of VLAN IDs. Note that you can not modify the VLAN name.
3. Click **Update**.
4. Click **Apply**.
5. At the top of the window, click **Save Configuration.**

Deleting a VLAN Pool

1. On the **VLAN Pool** window, click **Delete** next to the VLAN name you want to delete. A prompt appears.
2. Click **OK**.
3. Click **Apply**.
4. At the top of the window, click **Save Configuration**.

Creating a VLAN Pool Using the CLI



VLAN pooling should *not* be used with static IP addresses.

This example creates a VLAN pool named **mygroup** that has the assignment type **even**.

```
(host)(config) #vlan-name mygroup pool assignment even
```

Viewing and Adding VLAN IDs Using the CLI

The following example shows how to view VLAN IDs in a VLAN pool:

```
(host) #configure terminal
Enter Configuration commands, one per line. End with CNTL/Z
(host)(config) #show vlan

VLAN CONFIGURATION
------------------
VLAN   Description  Ports
----   -----------  -----
1      Default      FE1/0-3 FE1/6 GE1/8
2      VLAN0002
4      VLAN0004
12     VLAN0012
210    VLAN0210
212    VLAN0212     FE1/5
213    VLAN0213     FE1/4
1170   VLAN1170     FE1/7
1170   VLAN1170     FE1/7
```

The following example shows how to add existing VLAN IDs to a VLAN pool:

```
(host) #configure terminal
Enter Configuration commands, one per line. End with CNTL/Z
(host)(config) #vlan-name mygroup pool
(host)(config) #vlan mygroup 2,4,12
```

```
(host)(config) #
```

To confirm the VLAN pool status and mappings assignments, use the **show vlan mapping** command:

```
(host)(config) #show vlan mapping

Vlan Mapping Table
------------------
VLAN Name       Pool Status  Assignment Type  VLAN IDs
---------       -----------  ---------------  --------
mygroup         Enabled      Hash             62,94
newpoolgroup    Enabled      Even
vlannametest    Enabled      Even             62,1511
```

Role Derivation for Named VLAN Pools

You can configure Named VLANs under user rule, server derivation, user derivation, and VSA in this release.



.You cannot modify a VLAN name, so choose the name carefully.

Named VLANs (single VLAN IDs or VLAN pools) can only be assigned to tunnel mode VAP's and wired profiles. They can also be assigned to user roles, user rule derivation, server derivation, and VSA for tunnel and bridge mode.

For tunnel mode, VLAN pools that have the assignment type "hash" and "even" are supported.

For bridge mode only, VLAN pools with the assignment type "hash" are supported. If a VLAN pool with "even" assignment is assigned to a user rule, user role, server derivation or VSA, than the "hash" assignment is applied and the following error message displays:

"vlan pool assignment type EVEN not supported for bridge. Applying HASH algorithm to retrieve vlan-id"

Note that L2 roaming is not supported with an even VLAN assignment.

In the CLI

To apply a named VLAN pool name in a user rule, use the existing CLI commands:

```
(host)(config) #aaa derivation-rules
(host)(config) #aaa derivation-rules user <string>
(host)(config) #aaa derivation-rules user test-user-rule
(host)(user-rule) #set vlan
```

To apply a named VLAN pool in a user role, use the existing CLI commands:

```
(host)(config) #user-role test-vlan-name
(user)(config-role) #vlan test-vlan
```

To apply a named VLAN pool in server derivation, use the CLI commands:

```
(host)(config) #aaa server-group test-vlan-server-group
(user)(Server Group "test-vlan-server-group") set vlan
```

For a named VLAN derivation using VSA, configure the RADIUS server using these values:

```
Aruba-Named-UserVLAN   9   String   Aruba   14823
```

In the WebUI

To apply a named VLAN pool in a user rule, navigate to the WebUI page:

Security > Authentication > User Rules

To apply a named VLAN pool in a user role, navigate to the WebUI page:

Security > Access Control > User Roles > Add or Edit Role

To apply a named VLAN pool in a server derivation (server group), navigate to the WebUI page:

Security > Authentication> Servers > Server Group > <server-group_name> >Server Rules

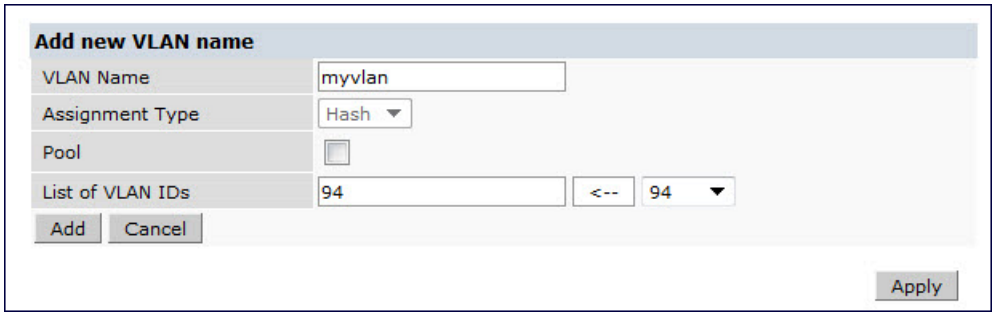
Creating a Named VLAN not in a Pool

The following configuration assigns the name **myvlan** to the VLAN ID 94:

In the WebUI

1. Navigate to **Configuration > Network > VLANs**.
2. Select the **VLAN Pool**tab to open the **VLAN Pool** window.
3. Click **Add**.
4. In the **VLAN Name** field, enter a name that identifies this VLAN.
5. Make sure the **Pool** field is unchecked. The **Assignment Type** is grayed out as this field applies only to VLAN pools.

Figure 15 *Named VLAN not in a Pool*



6. In the **List of VLAN IDs** field, enter the VLAN ID you want to name. If you know the ID, enter the ID. You can also click the drop-down list to view the IDs, then click the **<--** arrow to add the ID to the pool.
7. Click **Apply**.

In the CLI

This example assigns a name to an existing VLAN ID.

```
(host)(config) #vlan-name myvlan
(host)(config) #vlan myvlan 94
```

This example assigns a VLAN name in a virtual AP:

```
(host)(config) #wlan virtual-ap default vlan mygroup
```

This example assigns a VLAN name in a wired profile for access VLAN:

```
(host)(Wired AP profile "default") #switchport access vlan mygroup
```

This example assigns a VLAN name in a wired profile for a trunk VLAN and an allowed VLAN:

```
(host)(Wired AP profile "default") #switchport access vlan mygroup

(host)(config) #ap wired-ap-profile default switchport trunk ?
allowed                 Set allowed VLAN characteristics when interface is
                        in trunking mode
native                  Set trunking native characteristics when interface
                        is in trunking mode

(host)(config) #ap wired-ap-profile default switchport trunk native vlan mynativevlan
```

```
(host)(config) #ap wired-ap-profile default switchport trunk allowed vlan myallowedvlan
```

Adding a Bandwidth Contract to the VLAN

Bandwidth contracts on a VLAN can limit broadcast and multicast traffic. ArubaOS includes an internal exception list to allow broadcast and multicast traffic using the VRRP, LACP, OSPF, PVST, and STP protocols. To remove per-VLAN bandwidth contract limits on an additional broadcast or multicast protocol, add the MAC address for that broadcast/multicast protocol to the VLAN Bandwidth Contracts MAC Exception List.

The command in the example below adds the MAC address for CDP (Cisco Discovery Protocol) and VTP (Virtual Trunking Protocol to the list of protocols that are not limited by VLAN bandwidth contracts.

```
(host)(config) #vlan-bwcontract-explist mac 01:00:0C:CC:CC:CC
```

To show entries in the VLAN bandwidth contracts MAC exception list, use the **`show vlan-bwcontract-explist [internal]`** command:

```
(host)(config) #show vlan-bwcontract-explist internal

VLAN BW Contracts Internal MAC Exception List
---------------------------------------------
MAC address
-----------
01:80:C2:00:00:00
01:00:0C:CC:CC:CD
01:80:C2:00:00:02
01:00:5E:00:82:11
```

Optimizing VLAN Broadcast and Multicast Traffic

Broadcast and Multicast (BCMC) traffic from APs, remote APs, or distributions terminating on the same VLAN floods all VLAN member ports. This causes critical bandwidth wastage, especially when the APs are connected to an L3 cloud where the available bandwidth is limited or expensive. Suppressing the VLAN BCMC traffic to prevent flooding can result in loss of client connectivity.

To effectively prevent flooding of BCMC traffic on all VLAN member ports, use the **`bcmc-optimization`** parameter under the `interface vlan` command. This parameter ensures controlled flooding of BCMC traffic without compromising the client connectivity. This option is disabled by default. You must enable this parameter for the controlled flooding of BCMC traffic.



If you enable BCMC Optimization on uplink ports, the controller-generated Layer-2 packets will be dropped.

The **bcmc-optimization** parameter has the following exemptions:

- All DHCP traffic will continue to flood VLAN member ports even if you enable the **bcmc-optimization** parameter.
- ARP broadcasts and VRRP (multicast) traffic will still be allowed.

You can configure BCMC optimization using the CLI or WebUI.

Using the CLI

```
(host)(config) #interface vlan 1
(host)(config-subif)#bcmc-optimization
(host)(config-subif)#show interface vlan 1

VLAN1 is up line protocol is up
Hardware is CPU Interface, Interface address is 00:0B:86:61:5B:98 (bia 00:0B:86:61:5B:98)
Description: 802.1Q VLAN
```

```
Internet address is 10.17.22.1  255.255.255.0
Routing interface is enable, Forwarding mode is enable
Directed broadcast is disabled, BCMC Optimization enable
Encapsulation 802, loopback not set
MTU 1500 bytes
Last clearing of "show interface" counters 12 day 1 hr 4 min 12 sec
link status last changed 12 day 1 hr 2 min 21 sec
```

Proxy Arp is disabled for the Interface

Using the WebUI

1. Navigate to **Configuration > Network > IP**.
2. In the **IP Interfaces** tab, click **Edit** of the VLAN for configuring BCMC optimization.
3. Select the **Enable BCMC** check box to enable BCMC Optimization for the selected VLAN.

Figure 16 *Enable BCMC Optimization*



Configuring Ports

Both Fast Ethernet and Gigabit Ethernet ports can be set to access or trunk mode. A port is in access mode enabled by default and carries traffic only for the VLAN to which it is assigned. In trunk mode, a port can carry traffic for multiple VLANs.

For a trunk port, specify whether the port will carry traffic for all VLANs configured on the controller or for specific VLANs only. You can also specify the native VLAN for the port. A trunk port uses 802.1q tags to mark frames for specific VLANs, However, frames on a native VLAN are not tagged.

Classifying Traffic as Trusted or Untrusted

You can classify wired traffic based not only on the incoming physical port and channel configuration, but also on the VLAN associated with the port and channel.

About Trusted and Untrusted Physical Ports

Physical ports on the controller are trusted and usually connected to internal networks by default, while untrusted ports connect to third-party APs, public areas, or other networks to which you can apply access controls. When you define a physical port as untrusted, traffic passing through that port needs to go through a predefined access control list policy.

About Trusted and Untrusted VLANs

You can also classify traffic as trusted or untrusted based on the VLAN interface and port or channel. This means that wired traffic on the incoming port is trusted only when the port's associated VLAN is also trusted; otherwise the traffic is untrusted. When a port and its associated VLANs are untrusted, any incoming and outgoing traffic must pass through a predefined ACL. For example, this setup is useful if your company provides wired user guest access, and you want guest user traffic to pass through an ACL to connect to a captive portal.

You can set a range of VLANs as trusted or untrusted in trunk mode. The following table lists the port, VLAN and the trust/untrusted combination to determine if traffic is trusted or untrusted. Both the port and the

VLAN have to be configured as trusted for traffic to be considered as trusted. If the traffic is classified as untrusted, then traffic must pass through the selected session access control list and firewall policies.

Table 29: *Classifying Trusted and Untrusted Traffic*

Port	VLAN	Traffic Status
Trusted	Trusted	Trusted
Untrusted	Untrusted	Untrusted
Untrusted	Trusted	Untrusted
Trusted	Untrusted	Untrusted

Configuring Trusted/Untrusted Ports and VLANs

You can configure an Ethernet port as an untrusted access port, assign VLANs and classify them as untrusted, and designate a policy through which VLAN traffic on this port must pass.

In the WebUI

1. Navigate to the **Configuration > Network > Ports** window.
2. In the **Port Selection** section, click the port you want to configure.
3. In the **Make Port Trusted** section, clear the **Trusted** check box to make the port untrusted. The default is trusted (checked).
4. In the **Port Mode** section, select **Access**.
5. From the **VLAN ID** drop-down list, select the **VLAN ID** whose traffic will be carried by this port.
6. In the **Enter VLAN(s)** section**,** clear the **Trusted** check box to make the VLAN untrusted. The default is trusted (checked).
7. In the **VLAN Firewall Policy** drop-down list, select the policy through which VLAN traffic must pass. You can select a policy for both trusted and untrusted VLANs.
8. From the **Firewall Policy** section, select the policy from the **in** drop-down list through which inbound traffic on this port must pass.
9. Select the policy from the **out** drop-down list through which outbound traffic on this port must pass.
10. To apply a policy to this session's traffic on this port and VLAN, select the policy from the **session** drop-down list.
11. Click **Apply**.

In the CLI

In this example,

```
(host)(config) #interface range fastethernet 1/2
(host)(config-if)#switchport mode access
(host)(config-if)#no trusted
(host)(config-if)#switchport access vlan 2
(host)(config-if)#no trusted vlan 2
(host)(config-if)#ip access-group ap-acl session vlan 2
(host)(config-if)#ip access-group validuserethacl in
(host)(config-if)#ip access-group validuserethacl out
(host)(config-if)#ip access-group validuser session
```

Configuring Trusted and Untrusted Ports and VLANs in Trunk Mode

The following procedures configure a range of Ethernet ports as untrusted native trunks ports, assign VLANs and classify them as untrusted, and designate a policy through which VLAN traffic on the ports must pass.

In the WebUI

1. Navigate to the **Configuration > Network > Ports** window.
2. In the **Port Selection** section, click the port you want to configure.
3. For **Port Mode** select **Trunk**.
4. To specify the native VLAN, select a VLAN from the **Native VLAN** drop-down list and click the <-- arrow.
5. Choose one of the following options to control the type of traffic the port carries:
   - **Allow All VLANS Except**: The port carries traffic for all VLANs except those from this drop-down list.
   - **Allow VLANs**: The port carries traffic for all VLANs selected from this drop-down list.
   - **Remove VLANs:** The port does not carry traffic for any VLANs selected from this drop-down list.
6. To designate *untrusted* VLANs on this port, click **Trusted except**. In the corresponding VLAN field enter a range of VLANs that you want to make *untrusted*. (In this format, for example: 200-300, 401-500 and so on). Only VLANs listed in this range are untrusted. To designate only one VLAN as untrusted, select a VLAN from the drop-down list.
7. To designate *trusted* VLANs on this port, click **Untrusted except**. In the corresponding VLAN field, enter a range of VLANs that you want to designate as *trusted*. (In this format, for example: 200-300, 401-500 and so on). Only VLANs listed in this range are trusted. To designate only one VLAN as trusted, select a VLAN from the drop-down menu.
8. To remove a VLAN, click the **Remove VLANs** option and select the VLAN you want to remove from the drop-down list, and click the left arrow to add it back to the list.
9. To designate the policy through which VLAN traffic must pass, click **New** under the **Session Firewall Policy** field.
10. Enter the VLAN ID or select it from the associated drop-down list. Then select the policy, through which the VLAN traffic must pass, from the **Policy** drop-down list and click **Add**. Both the selected VLAN and the policy appear in the **Session Firewall Policy** field.
11. When you are finished listing VLANs and policies, click **Cancel**.
12. Click **Apply.**

In the CLI

Use the following examples:

```
(host)(config) #interface fastethernet 2/0
(host)(config-if)#description FE2/
(host)(config-if)#trusted vlan 1-99,101, 104, 106-199, 201-299
(host)(config-range)# switchport mode trunk
(host) (config-if)#switchport trunk native vlan 100
(host) (config-range)# ip access-group
(host) (config-range)# ip access-group test session vlan 2
```

Understanding VLAN Assignments

A client is assigned to a VLAN by one of several methods, in order of precedence. The assignment of VLANs are (from lowest to highest precedence):

1. The default VLAN is the VLAN configured for the WLAN (see Virtual AP Profiles on page 395).

2. Before client authentication, the VLAN can be derived from rules based on client attributes (SSID, BSSID, client MAC, location, and encryption type). A rule that derives a specific VLAN takes precedence over a rule that derives a user role that may have a VLAN configured for it.
3. After client authentication, the VLAN can be configured for a default role for an authentication method, such as 802.1x or VPN.
4. After client authentication, the VLAN can be derived from attributes returned by the authentication server (*server-derived rule*). A rule that derives a specific VLAN takes precedence over a rule that derives a user role that may have a VLAN configured for it.
5. After client authentication, the VLAN can be derived from Microsoft Tunnel attributes (Tunnel-Type, Tunnel Medium Type, and Tunnel Private Group ID). All three attributes must be present as shown below. This does not require a server-derived rule. For example:

   ```
   Tunnel-Type="VLAN"(13)
   Tunnel-Medium-Type="IEEE-802" (6)
   Tunnel-Private-Group-Id="101"
   ```

6. After client authentication, the VLAN can be derived from Vendor Specific Attributes (VSA) for RADIUS server authentication. This does not require a server-derived rule. If a VSA is present, it overrides any previous VLAN assignment. For example:

   ```
   Dell-User-VLAN
   Dell-Named-User-VLAN
   ```

VLAN Derivation Priorities for VLAN types

The VLAN derivation priorities for VLAN is defined below in the increasing order:

1. Default or Virtual AP VLAN
2. VLAN from Initial role
3. VLAN from User Derivation Rule (UDR) role
4. VLAN from UDR
5. VLAN from DHCP option 77 UDR role (wired clients)
6. VLAN from DHCP option 77 UDR (wired clients)
7. VLAN from MAC-based Authentication default role
8. VLAN from Server Derivation Rule (SDR) role during MAC-based Authentication
9. VLAN from SDR during MAC-based Authentication
10. VLAN from Vendor Specific Attributes (VSA) role during MAC-based Authentication
11. VLAN from VSA during MAC-based Authentication
12. VLAN from Microsoft Tunnel attributes during MAC-based Authentication
13. VLAN from 802.1X default role
14. VLAN from SDR role during 802.1X
15. VLAN from SDR during 802.1X
16. VLAN from VSA role during 802.1X
17. VLAN from VSA during 802.1X
18. VLAN from Microsoft Tunnel attributes during 802.1X
19. VLAN from DHCP options role
20. VLAN from DHCP options



A VLAN from DHCP options has highest priority for VLAN derivation. Note, however, that DHCP options are not considered for derivation if the Aruba VSA **ARUBA_NO_DHCP_FINGERPRINT (14)** was sent for the user.

Use the following command to display user VLAN derivation related debug information:

```
(host) #show aaa debug vlan user [ip | ipv6 | mac]
```

How a VLAN Obtains an IP Address

A VLAN on the controller obtains its IP address in one of the following ways:

- You can manually configure it. This is the default method and is described in Assigning a Static Address to a VLAN on page 159. At least one VLAN on the controller must be assigned a static IP address.
- Dynamically assigned from a Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE) server.

Assigning a Static Address to a VLAN

You can manually assign a static IP address to a VLAN on the controller. At least one VLAN on the controller a static IP address.

In the WebUI

1. Navigate to the **Configuration > Network > IP > IP Interfaces** page on the WebUI. Click **Edit** for the VLAN you just added.
2. Select the **Use the following IP address** option. Enter the IP address and network mask of the VLAN interface. If required, you can also configure the address of the DHCP server for the VLAN by clicking **Add**.
3. Click **Apply**.

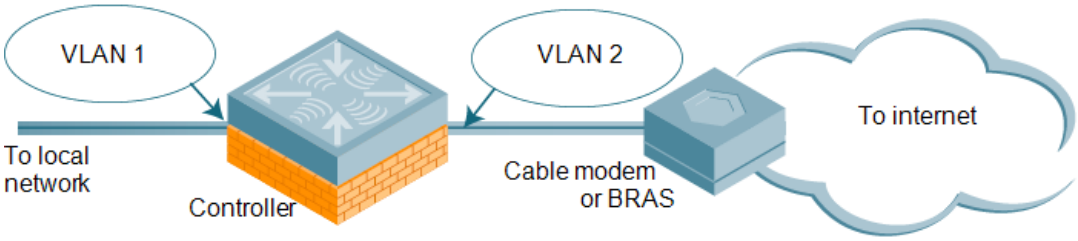
In the CLI

```
(host)(config) #interface vlan <id>
   ip address <address> <netmask>
```

Configuring a VLAN to Receive a Dynamic Address

In a branch office, you can connect a controller to an uplink switch or server that dynamically assigns IP addresses to connected devices. For example, you can connect the controller to a DSL or cable modem, or a broadband remote access server (BRAS). The following figure shows a branch office where a controller connects to a cable modem. VLAN 1 has a static IP address, while VLAN 2 has a dynamic IP address assigned via DHCP or PPPoE from the uplink device.

Figure 17 *IP Address Assignment to VLAN via DHCP or PPPoE*



Configuring Multiple Wired Uplink Interfaces (Active-Standby)

You can assign up to four VLAN interfaces to operate in active-standby topology. An active-standby topology provides redundancy so that when an active interface fails, the user traffic can failover to the standby interface.

To allow the controller to obtain a dynamic IP address for a VLAN, enable the DHCP or PPPoE client on the controller for the VLAN.

The following restrictions apply when enabling the DHCP or PPPoE client on the controller:

- You can enable the DHCP/PPPoE client multiple uplink VLAN interfaces (up to four) on the controller; these VLANs cannot be VLAN 1.
- Only one port in the VLAN can be connected to the modem or uplink switch.
- At least one interface in the VLAN must be in the up state before the DHCP/PPPoE client requests an IP address from the server.

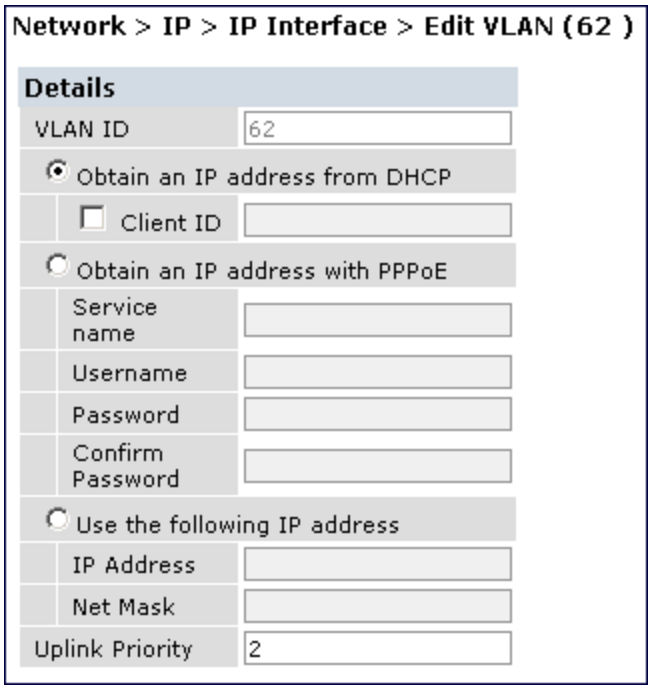
Enabling the DHCP Client

The DHCP server assigns an IP address for a specified amount of time called a lease. The controller automatically renews the lease before it expires. When you shut down the VLAN, the DHCP lease is released.

In the WebUI

1. Navigate to the **Configuration > Network > IP > IP Interfaces** page.
2. Click **Edit** for a previously-created VLAN.
3. Select **Obtain an IP address from DHCP**.
4. Enter a priority value for the VLAN ID in the **Uplink Priority** field. All wired uplink interfaces have the same priority by default. If you want to use an active-standby topology, then prioritize each uplink interfaces by entering a different priority value (1– 4) for each uplink interface.

Figure 18 *Assigning VLAN Uplink Priority—Active-Standby Configuration*



5. Click **Apply**.

In the CLI

In this example, the DHCP client has the client ID name *myclient*, and the interface VLAN 62 has an uplink priority of 2:

```
interface vlan 62
uplink wired vlan 62 priority 2
interface vlan 62 ip address dhcp-client client-id myclient
```

Enabling the PPPoE Client

To authenticate the BRAS and request a dynamic IP address, the controller must have the following configured:

- PPPoE user name and password to connect to the DSL network
- PPPoE service name: either an ISP name or a class of service configured on the PPPoE server

When you shut down the VLAN, the PPPoE session terminates.

In the WebUI

1. Navigate to the **Configuration > Network > IP > IP Interfaces** page.
2. Click **Edit** for a previously-created VLAN.
3. Select **Obtain an IP address with PPPoE.**
4. Enter the service name, username, and password for the PPPoE session.
5. Enter a priority value for the VLAN ID in the **Uplink Priority** field. All wired uplink interfaces have the same priority by default. If you want to use an active-standby topology, then prioritize each uplink interfaces by entering a different priority value (1– 4) for each uplink interface.
6. Click **Apply.**

In the CLI

In this example, a PPoE service name, username, and password are assigned, and the interface VLAN 14 has an uplink priority of 3:

```
(host)(config) #interface vlan 14
  ip address pppoe
(host)(config) #interface vlan 14 ip pppoe-service-name <service_name>
(host)(config) #interface vlan 14 ip pppoe-username <username>
(host)(config) #(host) (config) #interface vlan 14 ip pppoe-password *****
(host)(config) #uplink wired vlan 14 priority 3
```

Default Gateway from DHCP/PPPoE

You can specify that the router IP address obtained from the DHCP or PPPoE server be used as the default gateway for the controller.

In the WebUI

1. Navigate to the **Configuration > Network > IP > IP Routes** page.
2. For Default Gateway, select **(Obtain an IP address automatically)**.
3. Click **Apply**.

In the CLI

```
(host) (config) #ip default-gateway import
```

Configuring DNS/WINS Server from DHPC/PPPoE

The DHCP or PPPoE server can also provide the IP address of a DNS server or NetBIOS name server, which can be passed to wireless clients through the controller's internal DHCP server.

For example, the following configures the DHCP server on the controller to assign addresses to authenticated employees; the IP address of the DNS server obtained by the controller via DHCP/PPPoE is provided to clients along with their IP address.

In the WebUI

1. Navigate to the **Configuration > Network > IP > DHCP Server** page.

2. Select **Enable DCHP Server**.
3. Under Pool Configuration, select **Add**.
4. For Pool Name, enter employee-pool.
5. For Default Router, enter 10.1.1.254.
6. For DNS Servers, select **Import from DHCP/PPPoE**.
7. For WINS Servers, select **Import from DHCP/PPPoE**.
8. For Network, enter 10.1.1.0 for IP Address and 255.255.255.0 for Netmask.
9. Click **Done**.

In the CLI

Use the following commands:

```
(host)(config) #ip dhcp pool employee-pool
  d>efault-router 10.1.1.254
  d>ns-server import
  netbios-name-server import
  network 10.1.1.0 255.255.255.0
```

Configuring Source NAT to Dynamic VLAN Address

When a VLAN interface obtains an IP address through DHCP or PPPoE, a NAT pool (dynamic-srcnat) and a session ACL (dynamic-session-acl) are automatically created which reference the dynamically-assigned IP addresses. This allows you to configure policies that map private local addresses to the public address(es) provided to the DHCP or PPPoE client. Whenever the IP address on the VLAN changes, the dynamic NAT pool address also changes to match the new address.

For example, the following rules for a guest policy deny traffic to internal network addresses. Traffic to other (external) destinations are source NATed to the IP address of the DHCP/PPPoE client on the controller.

In the WebUI

1. Navigate to the **Configuration > Security > Access Control > Policies** page. Click **Add** to add the policy **guest**.
2. To add a rule, click **Add**.
   a. For Source, select **any**.
   b. For Destination, select **network** and enter **10.1.0.0** for Host IP and **255.255.0.0** for Mask.
   c. For Service, select **any**.
   d. For Action, select **reject**.
   e. Click **Add**.
3. To add another rule, click **Add.**
   a. Leave Source, Destination, and Service as **any.**
   b. For Action, select **src-nat**.
   c. For NAT Pool, select **dynamic-srcnat**.
   d. Click **Add**.
4. Click **Apply.**

In the CLI

Use the following commands:

```
(host)(config) #ip access-list session guest
  any network 10.1.0.0 255.255.0.0 any deny
  any any any src-nat pool dynamic-srcnat
```

Configuring Source NAT for VLAN Interfaces

The example configuration in the previous section illustrates how to configure source NAT using a policy that is applied to a user role. You can also enable source NAT for a VLAN interface to perform NAT on the source address for *all* traffic that exits the VLAN.

Packets that exit the VLAN are given a source IP address of the "outside" interface, which is determined by the following:

- If you configure "private" IP addresses for the VLAN, the controller is assumed to be the default gateway for the subnetwork. Packets that exit the VLAN are given the IP address of the controller for their source IP address.
- If the controller is forwarding the packets at Layer-3, packets that exit the VLAN are given the IP address of the next-hop VLAN for their source IP address.

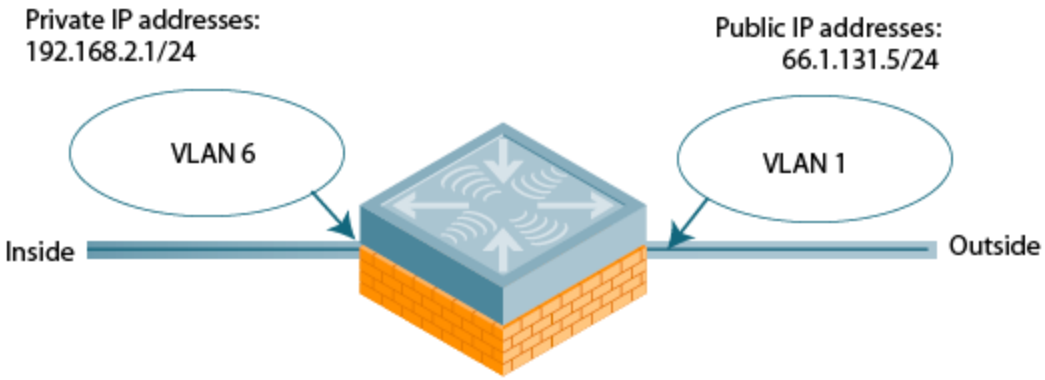


Do not enable the **NAT translation for inbound traffic** option for VLAN 1, as this will prevent IPsec connectivity between the controller and its IPsec peers.

Example Configuration

In the following example, the controller operates within an enterprise network. VLAN 1 is the outside VLAN, and traffic from VLAN 6 is source NATed using the IP address of the controller. The IP address assigned to VLAN 1 is used as the controller's IP address; thus traffic from VLAN 6 would be source NATed to 66.1.131.5:

Figure 19 *Example: Source NAT using Controller IP Address*



In the WebUI

1. Navigate to the **Configuration > Network > VLANs** page. Click **Add** to configure VLAN 6 (VLAN 1 is configured through the Initial Setup).
   a. Enter **6** for the VLAN ID.
   b. Click **Apply**.
2. Navigate to the **Configuration > Network > IP > IP Interfaces** page.
3. Click **Edit** for VLAN 6:
   a. Select **Use the following IP address**.
   b. Enter **192.168.2.1** for the IP Address and **255.255.255.0** for the Net Mask.
   c. Select the **Enable source NAT for this VLAN** checkbox.
4. Click **Apply**.

In the CLI

Use the following commands:

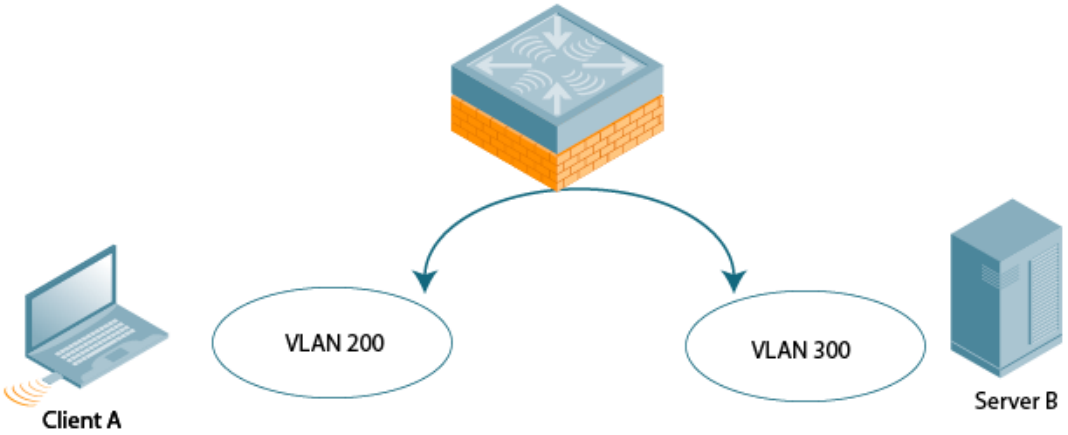
```
(host)(config) #interface vlan 1
   ip address 66.1.131.5 255.255.255.0
(host)(config) #interface vlan 6
(host)(config) #ip address 192.168.2.1 255.255.255.0
   ip nat inside
   ip default-gateway 66.1.131.1
```

Inter-VLAN Routing

On the controller, you can map a VLAN to a layer-3 subnetwork by assigning a static IP address and a netmask, or by configuring a DHCP or PPPoE server to provide a dynamic IP address and netmask to the VLAN interface. The controller, acting as a layer-3 switch, routes traffic between VLANs that are mapped to IP subnetworks; this forwarding is enabled by default.

In Figure 20, VLAN 200 and VLAN 300 are assigned the IP addresses 2.1.1.1/24 and 3.1.1.1/24, respectively. Client A in VLAN 200 is able to access server B in VLAN 300 and vice-versa, provided that there is no firewall rule configured on the controller to prevent the flow of traffic between the VLANs.

Figure 20 *Default Inter-VLAN Routing*



You can optionally disable layer-3 traffic forwarding to or from a specified VLAN. When you disable layer-3 forwarding on a VLAN, the following restrictions apply:

- Clients on the restricted VLAN can ping each other, but cannot ping the VLAN interface on the controller. Forwarding of inter-VLAN traffic is blocked.
- IP mobility does not work when a mobile client roams to the restricted VLAN. You must ensure that a mobile client on a restricted VLAN is not allowed to roam to a non-restricted VLAN. For example, a mobile client on a guest VLAN will not be able to roam to a corporate VLAN.

To disable layer-3 forwarding for a VLAN configured on the controller:

Using the WebUI to restrict VLAN routing

1. Navigate to the **Configuration > Network > IP > IP Interface** page.
2. Click **Edit** for the VLAN for which routing is to be restricted.
3. Configure the VLAN to either obtain an IP address dynamically (via DHCP or PPPoE) or to use a static IP address and netmask.
4. Deselect (uncheck) the **Enable Inter-VLAN Routing** checkbox.
5. Click **Apply**.

Using the CLI

Use the following commands:

```
interface vlan <id>
   ip address {<ipaddr> <netmask>|dhcp-client|pppoe}
   no ip routing
```

Configuring Static Routes

To configure a static route (such as a default route) on the controller, do the following:

In the WebUI

1. Navigate to the **Configuration > Network > IP > IP Routes** page.
2. Click **Add** to add a static route to a destination network or host. Enter the destination IP address and network mask (255.255.255.255 for a host route) and the next hop IP address.
3. Click **Done** to add the entry. Note that the route has not yet been added to the routing table.
4. Click **Apply** .. The message **Configuration Updated Successfully** confirms that the route has been added.

In the CLI

Use the following examples:

```
(host)(config) #ip route <address> <netmask> <next_hop>
```

Configuring the Loopback IP Address

The loopback IP address is a logical IP interface that is used by the controller to communicate with APs. The loopback address is used as the controller's IP address for terminating VPN and GRE tunnels, originating requests to RADIUS servers, and accepting administrative communications. You configure the loopback address as a host address with a 32-bit netmask. The loopback address is not bound to any specific interface and is operational at all times. To use this interface, ensure that the IP address is reachable through one of the VLAN interfaces. It will be routable from all external networks.

You must configure a loopback address if you are not using VLAN1 to connect the controller to the network. If you do not configure the loopback interface address, then the first configured VLAN interface address is selected. Generally, VLAN 1 is the factory default setting and thus becomes the controller IP address.

In the WebUI

1. Navigate to the **Configuration > Network > Controller > System Settings** page and locate the **Loopback Interface** section.
2. Modify the **IP Address** as required.
3. Click **Apply**.



If you are use the loopback IP address to access the WebUI, changing the loopback IP address will result in loss of connectivity. It is recommended that you use one of the VLAN interface IP addresses to access the WebUI.

4. Navigate to the **Maintenance > Controller > Reboot Controller** page to reboot the controller to apply the change of loopback IP address.
5. Click **Continue** to save the configuration.
6. When prompted that the changes were written successfully to flash, click **OK**.



7. The controller boots up with the changed loopback IP address.

In the CLI

Use the following commands:

```
(host)(config) #interface loopback ip address <address>
(host)(config) #write memory
```

Enter the following command in Enable mode to reboot the controller :

```
(host) #reload
```

Configuring the Controller IP Address

The Controller IP address is used by the controller to communicate with external devices such as APs.

NOTE

IP addresses used by the controller is not limited to the controller IP address.

You can set the Controller IP address to the loopback interface address or to an existing VLAN ID address. This allows you to force the controller IP address to be a specific VLAN interface or loopback address across multiple machine reboots. Once you configure an interface to be the controller IP address, that interface address cannot be deleted until you remove it from the controller IP configuration.

If the controller IP address is not configured then the controller IP defaults to the current loopback interface address. If the loopback interface address is not configured then the first configured VLAN interface address is selected. Generally, VLAN 1 is the factory default setting and thus becomes the controller IP address.

Using the WebUI:

1. Navigate to **Configuration > Network > Controller > System Settings** page.
2. Locate the Controller **IP Details** section.
3. Select the address you want to set the Controller IP to from the **VLAN ID** drop-down list. This list contains only VLAN IDs that have statically assigned IP addresses. If you have previously configured a loopback interface IP address, then it will also appear in this list. Dynamically assigned IP addresses such as DHCP/PPPOE do not display.
4. Click **Apply**.

NOTE

Any change in the controller's IP address requires a reboot.

5. Navigate to the **Maintenance > Controller > Reboot Controller** page to reboot the controller to apply the change of controller IP address.
6. Click **Continue** to save the configuration.
7. When prompted that the changes were written successfully to flash, click **OK**.



8. The controller boots up with the changed controller IP address. of the selected VLAN ID.

Using the CLI

```
(host)(config) #controller-ip [loopback|vlan <VLAN ID>]
```

Configuring GRE Tunnels

A controller supports generic routing encapsulation (GRE) tunnels between the controller and APs. An AP opens a GRE tunnel to the controller for each radio interface. On the AP, the other end of the GRE tunnel is specified by the IP address configured and the variable values (in descending order of priority): *<master>*, *<servername>*, and *<serverip>*. If these variable are left to default values, the AP uses DNS to look up Dell-master to discover the IP address of the controller.

The controller also supports GRE tunnels between the controller and other GRE-capable devices. Static IPv6 L2/L3 GRE tunnels can be established between Dell devices and other devices that support IPv6 GRE tunnel. IPv4 and IPv6 L2 GRE Tunnels carry both IPv6 and IPv4 traffic, and the IPv6 traffic can be redirected over the IPv4 L3 GRE Tunnel. This section describes how to configure a GRE tunnel to such a device, and how to direct traffic into the tunnel.



The controller uses GRE tunnels for communication between master and local controllers; these GRE tunnels are automatically created and are not subject to the configuration described in this section. This feature is supported only in IPv4.



If a VLAN interface has multiple IPv6 addresses configured, one of them is used as the tunnel source IPv6 address. If the selected IPv6 address is deleted from the VLAN interface, then the tunnel source IP is re-configured with the next available IPv6 address.

Important Points to Remember

- By default a GRE Tunnel Interface is in IPv4 L3 mode.
- IPv6 configurations will be allowed on an IPv4 Tunnel only if the tunnel mode is set to IPv6. Similarly, IPv4 configurations will be allowed on an IPv6 Tunnel only if the tunnel mode is set to IP.

Limitations

ArubaOS does not support the following functions for Static IPv6 GRE Tunnels:

- IPv6 Auto configuration and IPv6 Neighbor Discovery mechanisms do not apply to IPv6 tunnels.
- No support for Tunnel encapsulation limit and MTU discovery options on the IPv6 tunnels.
- You cannot use IPv6 GRE for a master-local setup as IPSec is not supported in this release.

Creating a Tunnel Interface

To create a GRE tunnel on the controller, you need to specify the following:

- Tunnel ID: a number between 1 and 2147483647.

- IP address and netmask for the tunnel.
- Tunnel source: the local endpoint for the tunnel on the controller. This can be one of the following:
  - Loopback address of the controller
  - A specified IP address
  - A specified VLAN
  - Controller-IP
  - IP Address
- Tunnel destination: the IP address of the remote endpoint of the tunnel on the other GRE device.

To create a L3 GRE tunnel on the controller, you need to specify the following:

- IP version : IPv4\IPv6
- Tunnel ID: this can be a number between 1 and 2147483647.
- Mode: Select L3
- IP Address and netmask for the tunnel.
- Tunnel source: the local endpoint for the tunnel on the controller. This can be one of the following:
  - Loopback address of the controller
  - A specified IP address
  - A specified VLAN
  - Controller-IP
  - IP Address

In the WebUI

1. Navigate to the **Configuration > Network > IP > GRE Tunnels** page.
2. Click **Add**.
3. Enter the tunnel ID.
4. Enter the IP address and netmask for the tunnel.
5. Select **Enabled** to enable the tunnel interface.
6. Select the tunnel source, if it is not the loopback address of the controller. If you select **IP Address**, enter the IP address for the tunnel source. If you select VLAN, select the ID of the VLAN.
7. Enter the IP address of the tunnel destination.
8. Click **Apply**.

In the CLI

Use the following commands:

```
(host)(config) #interface tunnel <number>
   description <string>
   inter-tunnel-flooding
   ip address <ipaddr> <ipmask>|internal
   ipv6 address <X:X:X:X::X>
   mtu <mtu>
   no ...
   shutdown
   trusted
   tunnel destination {<A.B.C.D>|ipv6 <X:X:X:X::X>}|keepalive <interval>|<retries>|mode gre
   {<num>|ip|ipv6}|source {<A.B.C.D>|controller-ip|ipv6 {X:X:X:X::X|controller-
   ip|loopback|vlan <vlan id>}|loopback|vlan <vlan id>}|vlan <vlan id>}
```

Directing Traffic into the Tunnel

You can direct traffic into the tunnel by configuring one of the following:

- Static route, which redirects traffic to the IP address of the tunnel



While redirecting traffic through L3 GRE Tunnel the controller's tunnel IP address should be used as the next-hop,instead of providing the destination IP address.

- Firewall policy (session-based ACL), which redirects traffic to the specified tunnel ID

Static Routes

You can configure a static route that specifies the IP address of a tunnel as the next-hop for traffic for a specific destination. See Configuring Static Routes on page 165 for descriptions of how to configure a static route.

Firewall Policy

You can configure a firewall policy rule to redirect selected traffic into a tunnel.

Traffic redirected by a firewall policy rule is *not* forwarded to a tunnel that is "down" (see Tunnel Keepalives on page 169 for more information on how GRE tunnel status is determined).

In the WebUI

1. Navigate to the **Configuration > Security > Access Control > Policies** page.
2. Click **Add** to create a new firewall policy, or click **Edit** to edit a specific policy.
3. Click **Add** to create a new policy rule.
4. Configure the Source, Destination, and Service for the rule.
5. For Action, select **redirect to tunnel**. Enter the tunnel ID.
6. Configure any additional options, and click **Add**.
7. Click **Apply**.

In the CLI

Use the following commands:

```
(host)(config) #ip access-list session <name>
  <source> <destination> <service> redirect tunnel <id>
```

Tunnel Keepalives

The controller can determine the status of a GRE tunnel by sending periodic keepalive frames on the L2 or L3 GRE tunnel. If you enable tunnel keepalives, the tunnel is considered "down" if there is repeated failure of the keepalives. If you configured a firewall policy rule to redirect traffic to the tunnel, traffic is not forwarded to the tunnel until it is "up". When the tunnel comes up or goes down, an SNMP trap and logging message is generated. The remote endpoint of the tunnel does not need to support the keepalive mechanism.

The controller sends keepalive frames at 60-second intervals by default and retries keepalives up to three times before the tunnel is considered down. You can reconfigure the intervals from the default. For the interval, specify a value between 1 and 86400 seconds. For the retries, specify a value between 0 and 1024.

In the WebUI

1. Navigate to the **Configuration > Network > IP > GRE Tunnels** page.
2. Click **Edit** for the tunnel for which you are enabling tunnel keepalives.
3. Select **Enable Heartbeats** to enable tunnel keepalives and display the Heartbeat Interval and Heartbeat Retries fields.
4. Enter values for **Heartbeat Interval** and **Heartbeat Retries**.

5. Click **Apply**.

In the CLI

Use the following commands:

```
(host)(config) #interface tunnel id
   tunnel keepalive [<interval> <retries>]
```

Configuring GRE Tunnel Group

ArubaOS provides redundancy for L3 generic routing encapsulation (GRE) tunnels. This feature enables automatic redirection of the user traffic to a standby tunnel when the primary tunnel goes down.

To enable this functionality, you must:

- configure a tunnel-group to group a set of tunnels.
- enable tunnel keepalives on all the tunnel interfaces assigned to the tunnel-group, and
- configure the session ACL with the tunnel-group as the redirect destination.



GRE Tunnel Redundancy is not applicable for GRE tunnels created for communications between controllers and APs.

Creating a Tunnel Group

A tunnel-group is identified by a name or number. You can add multiple tunnels to a tunnel-group. The order of the tunnels defined in the tunnel-group configuration specifies their standby precedence. The first member of the tunnel-group is the primary tunnel. When the first tunnel fails, the second tunnel carries the traffic. The third tunnel in the tunnel-group takes over if the second tunnel also fails. In the mean time, if the first tunnel comes up, it becomes the most eligible standby tunnel.



You can configure up to 32 tunnel-groups on a controller with a maximum of five tunnels in each tunnel-group.

You can also enable or disable pre-emption as part of the tunnel-group configuration. Pre-emption is enabled by default. The pre-emption option, automatically redirects the traffic whenever it detects an active tunnel with a higher precedence in the tunnel-group. When pre-emption is disabled, the traffic gets redirected to a higher precedence tunnel only when the tunnel carrying the traffic fails.

You can configure the tunnel-group using the WebUI or the CLI.

In the WebUI

1. Navigate to the **Configuration > Network > IP > GRE Tunnels** page.
2. Click **Add** under the **Tunnel Group** pane.
3. Specify a name for the tunnel-group in the **Tunnel Group Name** text box.
4. Specify the tunnel IDs with comma separators in the **Tunnel Group Member** text box.
5. Select the **Enable Preemptive-Failover Mode** check box to enable pre-emption (Default: enabled); clear the checkbox to disable pre-emption.
6. Click **Apply**.

In the CLI

Execute the following commands to configure a tunnel-group:

```
(host)(config) #tunnel-group <tunnel-group-name>
(host)(config-tunnel-group)# tunnel <tunnel-id>
```

Execute the following command to enable pre-emption:

```
(host)(config-tunnel-group)#preemptive-failover
```

Following is a sample configuration:

```
(host)(config) #tunnel-group tgroup1
(host)(config-tunnel-group)# tunnel 10
(host)(config-tunnel-group)# tunnel 20
(host)(config-tunnel-group)#preemptive-failover
```

Execute the following command to view the operational status of a tunnel-group and its members:

```
(host)(config-tunnel-group)#show tunnel-group tgroup1
Tunnel-Group Table Entries
--------------------------
Tunnel Group Tunnel Group Id Preemptive Failover Active Tunnel Id Tunnel Members
------------ --------------- -------------------- ---------------- --------------
tgroup1      16385           enabled                               10 20
```

Execute the following command to view the operational status of all the configured tunnel-groups:

```
(host)(config) #show tunnel-group
Tunnel-Group Table Entries
--------------------------
Tunnel Group Tunnel Group Id Preemptive Failover Active Tunnel Id Tunnel Members
------------ --------------- -------------------- ---------------- --------------
tgroup1      16385           enabled              0                10 20
tgroup2      16387           enabled              40               20 40 10
tgroup3      16386           enabled              0                20
```

Execute the following command to view the datapath Tunnel-Group table entries:

```
(host) #show datapath tunnel-group
Datapath Tunnel-Group Table Entries
-----------------------------------
Tunnel-Group Active Tunnel Members
------------ ------------- -------------------
16387        11            11
```

Jumbo Frame Support

Jumbo frames are the data frames that are larger than 1500 bytes and includes the Layer 2 header and Frame Check Sequence (FCS). Jumbo frames functionality can be configured on W-7200 Series controllers to support up to 9216 bytes of payload.

In centralized deployments, frames that are more than 1500 bytes in size are generated from AP to the controller during encryption and enabling AMSDU. Therefore, whenever the AP associates to the controller, jumbo frames are used to get the highest network performace. If this functionality is not supported, the data frames gets fragmented, which reduces the overall throughput of the network and makes the network slow.



ArubaOS supports jumbo frames between 11ac APs and W-7200 Series controllers only.

You can enable the jumbo frame support in the following scenarios:

- Tunnel node: In a tunneled node deployment, the wired clients connected on the tunneled nodes can send and receive the jumbo frames.
- L2/L3 GRE tunnels: When you establish a GRE tunnel between two controllers, the clients on one controller can send and receive jumbo frames from the clients on the other controller on enabling jumbo frames.
- Between wired clients: In a network where clients connect to the controller with jumbo frames enabled ports can send and receive the jumbo frames.

- Wi-Fi tunnel: A Wi-Fi tunnel can support an AMSDU jumbo frame for an AP (The maximum MTU supported is up to 9216 bytes).

Limitations for Jumbo Frame Support

This release of ArubaOS does not support the jumbo frames for the following scenarios:

- IPsec, IPIP, and xSec.
- IPv6 fragmentation/reassembly.

Configuring Jumbo Frame Support

You can use the WebUI or CLI to configure the jumbo frame support.

Using the WebUI

To enable jumbo frame support globally:

1. Navigate to the **Configuration > ADVANCED SERVICES > Stateful firewall > Global Setting** page.
2. Select the **Jumbo frames processing** checkbox to enable the jumbo frames support.
3. Enter the value of the MTU in the **Jumbo MTU [1789-9216] bytes** textbox.
4. Click **Apply**.

To enable jumbo frame support on a port:

1. Navigate to **Configuration > NETWORK > Ports** page.
2. Select the **Enable Jumbo MTU** checkbox to enable the jumbo frames support.
3. Click **Apply**.

To enable jumbo frame support on a port channel:

1. Navigate to the **Configuration > NETWORK > Port-Channel** page.
2. Select the **Enable Jumbo MTU checkbox** to enable the jumbo frames support.
3. Click **Apply**.

Using the CLI

To enable the jumbo frame support globally and to configure the MTU value:

```
(host)(config)# firewall jumbo mtu <val>
```

You can configure the MTU value between 1,789-9,216. The default MTU value is 9,216.

To disable the jumbo frame support:

```
(host)(config)# no firewall enable-jumbo-frames
```

In this case, the MTU value is considered as 9,216 (default).

To enable jumbo frame support on a port channel:

```
(host)(config)# interface port-channel <id> jumbo
```

To disable jumbo frame support on a port channel:

```
(host)(config)# interface port-channel <id> no jumbo
```

To enable jumbo frame support on a port:

```
(host)(config) # interface gigabitethernet <slot>/<module>/<port> jumbo
```

To disable jumbo frame support on a port:

```
(host)(config) # interface gigabitethernet <slot>/<module>/<port> no jumbo
```

Viewing the Jumbo Frame Support Status

Execute the following command to view the global status of the jumbo frame support:

```
(host)#show firewall
Global firewall policies
-------------------------
Policy                                      Action    Rate        Port
------                                      ------    ----        ----
Enforce TCP handshake before allowing data  Disabled
Prohibit RST replay attack                  Disabled
Deny all IP fragments                       Disabled
Prohibit IP Spoofing                        Enabled
Monitor ping attack                         Disabled
Monitor TCP SYN attack                      Disabled
Monitor IP sessions attack                  Disabled
Deny inter user bridging                    Disabled
Log all received ICMP errors                Disabled
Per-packet logging                          Disabled
Blacklist Grat ARP attack client            Disabled
Session mirror destination                  Disabled
Stateful SIP Processing                     Enabled
Allow tri-session with DNAT                 Disabled
Disable FTP server                          No
Blacklist ARP attack client                 Disabled
Monitor ARP attack                          Disabled
Monitor Gratuitous ARP attack               Enabled
50/sec
GRE call id processing                      Disabled
Session Idle Timeout                        Disabled
Broadcast-filter ARP                        Disabled
WMM content enforcement                     Disabled
Session VOIP Timeout                        Disabled
Stateful H.323 Processing                   Enabled
Stateful SCCP Processing                    Enabled
Only allow local subnets in user table      Disabled
Monitor/police CP attacks                   Disabled
Rate limit CP untrusted ucast traffic       Enabled   9765 pps
Rate limit CP untrusted mcast traffic       Enabled   1953 pps
Rate limit CP trusted ucast traffic         Enabled   65535 pps
Rate limit CP trusted mcast traffic         Enabled   1953 pps
Rate limit CP route traffic                 Enabled   976 pps
Rate limit CP session mirror traffic        Enabled   976 pps
Rate limit CP auth process traffic          Enabled   976 pps
Deny inter user traffic                     Disabled
Prohibit ARP Spoofing                       Disabled
Stateful VOCERA Processing                  Enabled
Stateful UA Processing                      Enabled
Stall Detection                             Disabled
Enforce bw contracts for broadcast traffic  Disabled
Multicast automatic shaping                 Disabled
Enforce TCP Sequence numbers                Disabled
AMSDU                                       Enabled
Jumbo Frames                                Enabled   MTU = 9216
Session-tunnel FIB                          Enabled
Prevent DHCP exhaustion                     Disabled
Stateful SIPS Processing                    Enabled
Deny source routing                         Disabled
Immediate Freeback                          Disabled
Session mirror IPSEC                        Disabled
```

Execute the following command to view the jumbo frame status on a port:

```
(host)# show interface gigabitethernet <slot>/<port>/<module>
```

```
Example:
(host)# show interface gigabitethernet 0/0/0
GE 0/0/0 is up, line protocol is up
Hardware is Gigabit Ethernet, address is 00:1A:1E:00:0D:09 (bia 00:1A:1E:00:0D:09)
Description: GE0/0/0 (RJ45 Connector)
Encapsulation ARPA, loopback not set
Configured: Duplex ( AUTO ), speed ( AUTO )
Negotiated: Duplex (Full), speed (1000 Mbps)
Jumbo Support is enabled on this interface MTU 9216
Last clearing of "show interface" counters 1 day 20 hr 32 min 38 sec
link status last changed 1 day 19 hr 37 min 57 sec
120719 packets input, 24577381 bytes
Received 84208 broadcasts, 0 runts, 0 giants, 780 throttles
0 input error bytes, 0 CRC, 0 frame
32939 multicast, 36511 unicast
19865402 packets output, 4953350248 bytes
0 output errors bytes, 0 deferred
0 collisions, 0 late collisions, 0 throttles
This port is TRUSTED
```

Execute the following command to view the jumbo frame status on a port channel:

```
(host)#show interface port-channel  <id>
Example:
(host)#show interface port-channel 6
Port-Channel 6 is administratively up
Hardware is Port-Channel, address is 00:1A:1E:00:0D:08 (bia 00:1A:1E:00:0D:08)
Description: Link Aggregate (LACP)
Spanning Tree is forwarding
Switchport priority: 0
Jumbo Support is enabled on this interface MTU 9216
Member port:
GE 0/0/4, Admin is up, line protocol is up
GE 0/0/5, Admin is up, line protocol is up
Last clearing of "show interface" counters 1 day 20 hr 32 min 43 sec
link status last changed 1 day 20 hr 29 min 58 sec
69425936 packets input, 15102169223 bytes
Received 27578 broadcasts, 0 runts, 0 giants, 0 throttles
0 input error bytes, 0 CRC, 0 frame
27568 multicast, 69398358 unicast
270782 packets output, 37271325 bytes
0 output errors bytes, 0 deferred
0 collisions, 0 late collisions, 0 throttles
Port-Channel 6 is TRUSTED
```

This chapter describes ArubaOS support for IPv6 features:

- Understanding IPv6 Notation on page 175
- Understanding IPv6 Topology on page 175
- Enabling IPv6 on page 176
- Enabling IPv6 Support for Controller and APs on page 176
- Filtering an IPv6 Extension Header (EH) on page 184
- Configuring a Captive Portal over IPv6 on page 184
- Working with IPv6 Router Advertisements (RAs) on page 184
- RADIUS Over IPv6 on page 188
- TACACS Over IPv6 on page 189
- DHCPv6 Server on page 190
- Understanding ArubaOS Supported Network Configuration for IPv6 Clients on page 193
- Managing IPv6 User Addresses on page 198
- Understanding IPv6 Exceptions and Best Practices on page 199

Understanding IPv6 Notation

The IPv6 protocol is the next generation of large-scale IP networks, it supports addresses that are 128 bits long. This allows 2^{128} possible addresses (versus 2^{32} possible IPv4 addresses).

Typically, the IP address assigned on an IPv6 host consists of a 64-bit subnet identifier and a 64-bit interface identifier. IPv6 addresses are represented as eight colon-separated fields of up to four hexadecimal digits each. The following are examples of IPv6 addresses:

```
2001:0000:0eab:DEAD:0000:00A0:ABCD:004E
```

The use of the "::" symbol is a special syntax that you can use to compress one or more group of zeros or to compress leading or trailing zeros in an address. The "::" can appear only once in an address.

For example, the address, `2001:0000:0dea:C1AB:0000:00D0:ABCD:004E` can also be represented as:

```
2001:0:eab:DEAD:0:A0:ABCD:4E – leading zeros can be omitted
2001:0:0eab:dead:0:a0:abcd:4e – not case sensitive
2001:0:0eab:dead::a0:abcd:4e - valid
2001::eab:dead::a0:abcd:4e - Invalid
```

IPv6 uses a "/" notation which describes the no: of bits in netmask, similar to IPv4.

```
2001:eab::1/128 – Single Host
2001:eab::/64 – Network
```

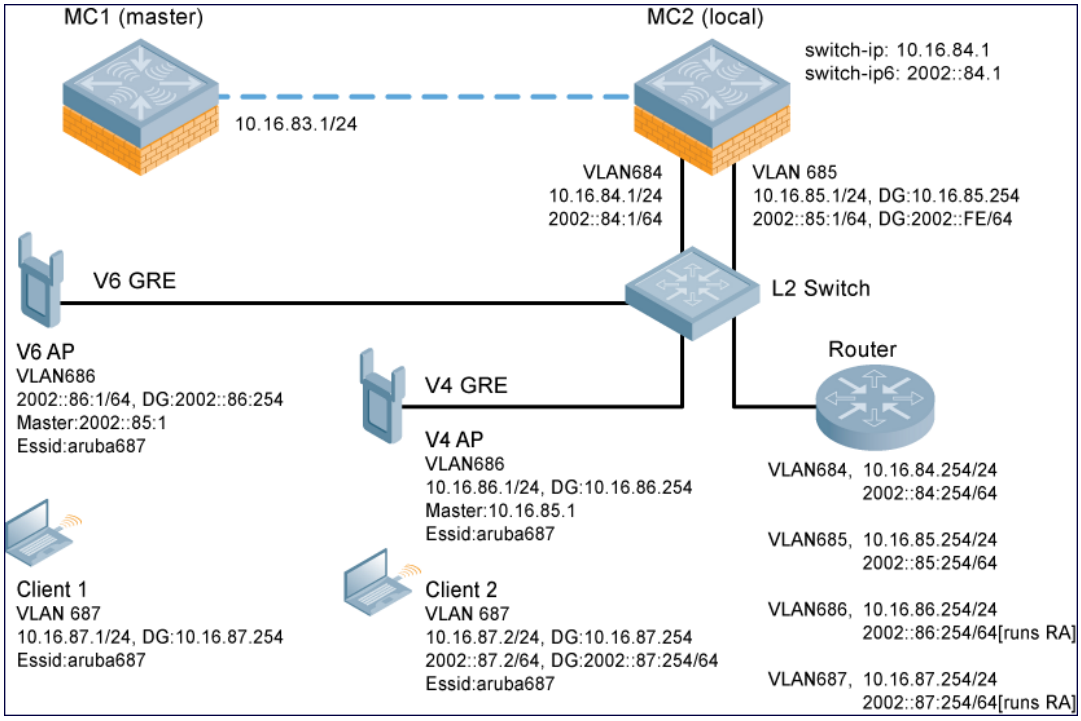
Understanding IPv6 Topology

IPv6 APs connect to the IPv6 controller over an IPv6 L3 network. The IPv6 controller can terminate both IPv4 and IPv6 APs. IPv4 and IPv6 clients can terminate to either IPv4 or IPv6 APs. ArubaOS supports Router Advertisements (RA). You do not need an external IPv6 router in the subnet to generate RA for IPv6 APs and clients that depend on stateless autoconfiguration to obtain IPv6 address. The external IPv6 router is the

default gateway in most deployments. However, the controller can be the default gateway by using static routes. The master-local communication always occurs in IPv4.

The following image illustrates how IPv6 clients, APs, and controllers communicate with each other in an IPv6 network:

Figure 21 *IPv6 Topology*



- The IPv6 controller (**MC2**) terminates both **V4 AP** (IPv4 AP) and **V6 AP** (IPv6 AP).
- **Client 1** (IPv4 client) terminates to **V6 AP** and **Client 2** (IPv6 client) terminates to **V4 AP**.
- **Router** is an external IPv6 router in the subnet that acts as the default gateway in this illustration.
- **MC1** (master) and **MC2** (local) communicates in IPv4.

Enabling IPv6

You must enable the IPv6 option on the controller before using any of the IPv6 functions. You can use the `ipv6 enable` command to enable the IPv6 packet/firewall processing on the controller. The IPv6 option is disabled by default.

You can also use the WebUI to enable the IPv6 option:

1. Navigate to the **Configuration > Advanced Services > Stateful Firewall** page.
2. Select the **Global Settings** tab.
3. Select the **IPv6 Enable** check box to enable the IPv6 option.
4. Click **Apply** .

Enabling IPv6 Support for Controller and APs

This release of ArubaOS provides IPv6 support for controllers and access points. You can now configure the master controller with an IPv6 address to manage the controllers and APs. Both IPv4 and IPv6 APs can

terminate on the IPv6 controller. You can provision an IPv6 AP in the network only if the controller interface is configured with an IPv6 address. An IPv6 AP can serve both IPv4 and IPv6 clients.



You must manually configure an IPv6 address on the controller interface to enable IPv6 support.

You can perform the following IPv6 operations on the controller:

- Configuring IPv6 Addresses on page 178
- Configuring IPv6 Static Neighbors on page 179
- Configuring IPv6 Default Gateway and Static IPv6 Routes on page 180
- Managing Controller IP Addresses on page 180
- Configuring Multicast Listener Discovery (MLD) on page 181
- Debugging an IPv6 Controller on page 183
- Provisioning an IPv6 AP on page 183

You can also view the IPv6 statistics on the controller using the following commands:

- `show datapath ip-reassembly ipv6` — View the IPv6 contents of the IP Reassembly statistics table.
- `show datapath route ipv6` — View datapath IPv6 routing table.
- `show datapath route-cache ipv6` — View datapath IPv6 route cache.
- `show datapath tunnel ipv6` — View the tcp tunnel table filtered on IPv6 entries.
- `show datapath user ipv6` — View datapath IPv6 user statistics such as current entries, pending deletes, high water mark, maximum entries, total entries, allocation failures, invalid users, and maximum link length.
- `show datapath session ipv6` — View datapath IPv6 session entries and statistics such as current entries, pending deletes, high water mark, maximum entries, total entries, allocation failures, invalid users, and maximum link length.

Additionally, you can view the IPv6 AP information on the controller using the following show commands:

- `show ap database`
- `show ap active`
- `show user`
- `show ap details ip6-addr`
- `show ap debug`

The following table lists IPv6 features:

Table 30: *IPv6 APs Support Matrix*

Features	Supported on IPv6 APs?
Forward Mode - Tunnel	Yes
Forward Mode - Decrypt Tunnel	No
Forward Mode - Bridge	No
Forward Mode - Split Tunnel	No
AP Type - CAP	Yes
AP Type - RAP	No

Features	Supported on IPv6 APs?
AP Type - Mesh Node	No
IPSEC	No
CPSec	No
Wired-AP/Secure-Jack	No
Fragmentation/Reassembly	Yes
MTU Discovery	Yes
Provisioning through Static IPv6 Addresses	Yes
Provisioning through IPv6 FQDN Master Name	Yes
Provisioning from WebUI	Yes
AP boot by Flash	Yes
AP boot by TFTP	No
WMM QoS	No
AP Debug and Syslog	Yes
ARM & AM	Yes
WIDS	Yes (Limited)
CLI support for users & datapath	Yes

Configuring IPv6 Addresses

You can configure IPv6 addresses for the management interface, VLAN interface, and the loopback interface of the controller. The controller can have up to three IPv6 addresses for each VLAN interface. The IPv6 address configured on the loopback interface or the first VLAN interface of the controller becomes the default IPv6 address of the controller.



If only one IPv6 address is configured on the controller, it becomes the default IPv6 address of the controller. With this release of ArubaOS, you can delete this IPv6 address.

You can configure IPv6 interface address using the WebUI or CLI. As per Internet Assigned Numbers Authority (IANA), Dell controllers support the following ranges of IPv6 addresses:

- Global unicast—2000::/3
- Unique local unicast—fc00::/7
- Link local unicast—fe80::/10

In the WebUI

To Configure Link Local Address

1. Navigate to the **Configuration > Network > IP** page and select the **IP Interfaces** tab.

2. Edit a VLAN # and select **IP version** as IPv6.
3. Enter the link local address in the **Link Local Address** field.
4. Click **Apply**.

To Configure Global Unicast Address

1. Navigate to the **Configuration > Network > IP** page and select the **IP Interfaces** tab.
2. Edit a VLAN # and select **IP version** as IPv6.
3. Enter the global unicast address and the prefix-length in the IP Address/Prefix-length field.
4. (Optional) Select the **EUI64 Format** check box, if applicable.
5. Click **Add** to add the address to the global address list.
6. Click **Apply**.

To Configure Loopback Interface Address

1. Navigate to the **Configuration > Network > Controller** page and select the **System Settings** tab.
2. Under **Loopback Interface** enter the loopback address in the **IPv6 Address** field.
3. Click **Apply**.



You cannot configure the management interface address using the WebUI.

In the CLI

To configure the link local address:

```
(host)(config)#interface vlan <vlan#>
(host)(config-subif)#ipv6 address <ipv6-address> link-local
```

To configure the global unicast address:

```
(host)(config)#interface vlan <vlan#>
(host)(config-subif)#ipv6 address <ipv6-prefix>/<prefix-length>
```

To configure the global unicast address (EUI 64 format):

```
(host)(config)#interface vlan <vlan#>
(host)(config-subif)#ipv6 address <ipv6-prefix/prefix-length> eui-64
```

To configure the management interface address:

```
(host)(config)#interface mgmt
(host)(config-subif)#ipv6 address <ipv6-prefix/prefix-length>
```

To configure the loopback interface address:

```
(host)(config)#interface loopback
(host)(config-subif)#ipv6 address <ipv6-prefix>
```

Configuring IPv6 Static Neighbors

You can configure a static neighbor on a VLAN interface either using the WebUI or the CLI.

In the WebUI

1. Navigate to the **Configuration > Network > IP** page and select the **IPv6 Neighbors** tab.
2. Click **Add** and enter the following details of the IPv6 neighbor:
   - IPV6 Address
   - Link-layer Addr
   - VLAN Interface

3. Click **Done** to apply the configuration.

In the CLI

To configure a static neighbor on a VLAN interface:

```
(host)(config)#ipv6 neighbor <ipv6addr> vlan <vlan#> <mac>
```

Configuring IPv6 Default Gateway and Static IPv6 Routes

You can configure IPv6 default gateway and static IPv6 routes using the WebUI or CLI.

In the WebUI

To Configure IPv6 Default Gateway

1. Navigate to the **Configuration > Network > IP** page and select the **IP Routes** tab.
2. Under the **Default Gateway** section, click **Add**.
3. Select IPv6 as **IP Version**, and enter the IPv6 address in the **IP Address** field.
4. Click **Add** to add the address to the IPv6 default gateway table.
5. Click **Apply**.

To Configure Static IPv6 Routes

1. Under the **IP Routes** section, click **Add** and select IPv6 as **IP Version**.
2. Enter the destination IP address and the forwarding settings in the respective fields.
3. Click **Done** to add the static route to the IPv6 routes table.
4. Click **Apply**.

In the CLI

To configure the IPv6 default gateway:

```
(host)(config)#ipv6 default-gateway <ipv6-address> <cost>
```

To configure static IPv6 routes:

```
(host)(config)#ipv6 route <ipv6-prefix/prefix-length> <ipv6-next-hop> <cost>
<ipv6-next-hop>  = X:X:X:X::X
```

Managing Controller IP Addresses

You can change the default controller IP address by assigning a different VLAN interface address or the loop back interface address. You can also turn on Syslog messaging for IPv6 (similar to IPv4 logging) using the `logging <ipv6 address>` command. For more information on logging, see Configuring Logging on page 804.You can use the WebUI or CLI to change the default controller IP address.

In the WebUI

1. Navigate to the **Configuration > Network > Controller** page and select the **System Settings** tab.
2. Under the **Controller IP Details** section, select the VLAN Id or the loopback interface Id in the **IPv6 Address** drop down.
3. Click **Apply.**

In the CLI

To configure an IPv6 address to the controller:

```
(host)(config)#controller-ipv6 loopback
(host)(config)#controller-ipv6 vlan <vlanId>
```

To enable logging over IPv6:

```
(host)(config)#logging <ipv6 address>
```

Configuring Multicast Listener Discovery (MLD)

You can enable the IPv6 multicast snooping on the controller by using the WebUI or CLI and configure MLD parameters such as query interval, query response interval, robustness variable, and ssm-range.

The Source Specific Multicast (SSM) supports delivery of multicast packets that originate only from a specific source address requested by the receiver. You can forward multicast streams to the clients if the source and group match the client subscribed source group pairs (S,G).

The controller supports the following IPv6 multicast source filtering modes:

- Include - In Include mode, the reception of packets sent to a specified multicast address is enabled only from the source addresses listed in the source list. The default IPv6 SSM address range is FF3X::4000:1 – FF3X::FFFF:FFFF, and the hosts subscribing to SSM groups can only be in the Include mode.
- Exclude - In Exclude mode, the reception of packets sent to a specific multicast address is enabled from all source addresses. If there is a client in the Exclude mode, the subscription is treated as an MLDv1 join.

For more information on MLD feature, see **RFC 3810** and **RFC 4604**,

In the WebUI

To enable IPv6 MLD Snooping

1. Navigate to the **Configuration > Network > IP** page and select the **IP Interfaces** tab.
2. Click the **Edit** button listed under **Actions** to edit the required VLAN interface.
3. Select IPv6 from the **IP version** drop-down list.
4. Check the **Enable MLD Snooping** check box under **MLD** section to enable IPv6 MLD snooping.
5. Click **Apply.**

To Modify IPv6 MLD Parameters

1. Navigate to the **Configuration > Network > IP** page and select the **Multicast** tab.
2. Under the **MLD** section, enter the required values in the following fields:
   - Robustness Variable: default value is 2
   - Query Interval (second): default value is 125 seconds
   - Query Response Interval (in 1/10 second): default value is 100 (1/10 seconds).
3. Click **Apply.**

To configure the SSM Range:

1. Navigate to **Configuration>Network>IP** page and select the **Multicast** tab.
2. In the **MLD** section, use the **SSM Range Start-IP** and **SSM Range End-IP** fields to configure the SSM Range.
3. Click **Apply** to save your changes.

In the CLI

To enable IPv6 MLD snooping:

```
(host)(config) #interface vlan 1
(host)(config-subif)#ipv6 mld snooping
```

To view if IPv6 MLD snooping is enabled:

```
(host)(config-subif)#show ipv6 mld interface
```

To view the MLD Group information:

```
(host)(config) #show ipv6 mld group
```

To modify IPv6 MLD parameters:

```
(host)(config) #ipv6 mld
(host)(config-mld) # query-interval <time in seconds (1-65535)>|query-response-interval <time in 1/10th of seconds (1-65535)|robustness-variable <value (2-10)>
```

To view MLD configuration:

```
(host)(config-subif)#show ipv6 mld config
```



When you enter the SSM Range ensure that the upstream router has the same range, else the multicast stream would be dropped.

Dynamic Multicast Optimization

When multiple clients are associated to an AP and when one client is subscribed for a multicast stream, all the clients associated to the AP receive the stream, as the packets are directed to the multicast MAC address. To restrict the multicast stream to only the subscribed clients, DMO sends the stream to the unicast MAC address of the subscribed clients. DMO is currently supported for both IPv4 and IPv6.

In the WebUI

You can configure the IPv6 DMO feature using the WebUI or CLI.

Using the WEBUI

To enable this feature using the WebUI:

1. Navigate to **Configuration>Wireless>AP Configuration** page.
2. Select the **AP Group** tab, click the AP Group you want to edit.
3. Expand the **Wireless LAN** menu, then expand the **Virtual AP** menu.
4. Select the Virtual AP profile for which you want to configure the Dynamic Multicast Optimization.
5. In the **Basic** tab under **Broadcast/Multicast** section configure the following parameters to enable multicasting:
   a. Select the **Dynamic Multicast Optimization (DMO)** checkbox,
   b. Use the **Dynamic Multicast Optimization (DMO) Threshold** field to set the maximum number of high-throughput stations in a multicast group.
6. Click **Apply** to save your changes.

In the CLI

To verify the DMO configuration, execute the following command:

```
(host) #show wlan virtual-ap
```

Limitations

The following are the MLDv2 limitations:

- Controller cannot route multicast packets.
- For mobility clients mld proxy should be used.
- VLAN pool scenario stream is forwarded to clients in both the VLANs even if the client from one of the VLANs is subscribed.
- Dynamic Multicast Optimization is applicable for wired clients in controllers.

Debugging an IPv6 Controller

ArubaOS provides the following debug commands for IPv6:

- `show ipv6 global` — displays if IPv6 is enabled globally or not
- `show ipv6 interface` — displays the configured IPv6 address, and any duplicate addresses
- `show ipv6 route/show datapath route ipv6` — displays the IPv6 routing information
- `show ipv6 ra status` — displays the Router Advertisement status
- `show Datapath session ipv6` — displays the IPv6 sessions created, and the sessions that are allowed
- `show datapath frame` — displays the IPv6 specific counters

You can also use the debug options such as ping and tracepath for IPv6 hosts. You can either use the WebUI or the CLI to use the ping and tracepath options.

In the WebUI

1. To ping an IPv6 host, navigate to the **Diagnostics > Network > Ping** page, enter an IPv6 address, and click **Ping**.
2. To trace the path of an IPv6 host, navigate to the **Diagnostics > Network > Tracepath** page, enter an IPv6 address, and click **Trace**.

In the CLI

To ping an IPv6 host:

```
(host)#ping ipv6 <global-ipv6-address>
(host)#ping ipv6 interface vlan <vlan-id> <linklocal-address>
```

To trace the path of an IPv6 host:

```
(host)#tracepath <global-ipv6-address>
```

Provisioning an IPv6 AP

You can provision an IPv6 AP on an IPv6 controller. You can either configure a static IP address or obtain a dynamic IPv6 address via stateless-autoconfig. The controller can act as the default gateway for the IPv6 clients, if static IPv6 routes are set on the controller.



Starting from ArubaOS 6.3, a wired client can connect to the Ethernet interface of an IPv6 enabled AP.

You can provision an IPv6 AP using the WebUI or CLI.

In the WebUI

1. Navigate to the **Configuration > AP Installation> Provision** page and select the **Provisioning** tab.
2. Select an AP and click **Provision**.
3. Under the **Master Discovery** section, enter the host controller IP address and the IPv6 address of the master controller.
4. To provision a static IP, select the **Use the following IP address** check box under the **IP Settings** section, and enter the following details:
   - IPv6 Address/Prefix-lengths
   - Gateway IPv6 Address
   - DNS IPv6 Address



Ensure that CPSEC is disabled before rebooting the AP.

5. Click **Apply and Reboot** to bring the IPv6 AP up.

In the CLI

To provision a static IPv6 address:

```
(host)(config)# provision-ap
```

Enhancements to IPv6 Support on AP

This release of ArubaOS provides the following IPv6 enhancements on the AP:

- DNS based ipv6 controller discovery
- FTP support for image upgrade in an IPv6 network
- DHCPv6 client support

Filtering an IPv6 Extension Header (EH)

ArubaOS firewall is enhanced to process the IPv6 Extension Header (EH) to enable IPv6 packet filtering. You can now filter the incoming IPv6 packets based on the EH type. You can edit the packet filter options in the default EH, using the CLI. The default EH alias permits all EH types.

Execute the following commands to permit or deny the IPv6 packets matching an EH type:

```
(host)(config) #netexthdr default
(host)(config-exthdr) #eh <eh-type> permit | deny
```

To view the EH types denied:

```
(host) (config-exthdr) #show netexthdr default

Extended Header type(s) Denied
------------------------------
51,  234,
```

Configuring a Captive Portal over IPv6

IPv6 is now enabled on the captive portal for user authentication on the Dell controller. For user authentication, use the internal captive portal that is initiated from the controller. A new parameter `captive` has been added to the IPv6 captive portal session ACL:

```
ipv6 user alias controller 6 svc-https captive
```



This release does not support external captive portal for IPv6. The captive portal authentication, customization of pages, and other attributes are same as IPv4.

You can configure captive portal over IPv6 (similar to IPv4) using the WebUI or CLI. For more information on configuration, see Configuring Captive Portal in the Base Operating System on page 300.

Working with IPv6 Router Advertisements (RAs)

ArubaOS enables the controllers to send router advertisements (RA) in an IPv6 network. Each host auto generates a link local address when you enable ipv6 on the host. The link local address allows the host to

communicate between the nodes attached to the same link.

The IPv6 stateless autoconfiguration mechanism allows the host to generate its own addresses using a combination of locally available information and information advertised by the routers. The host sends a router solicitation multicast request for its configuration parameters in the IPv6 network. The source address of the router solicitation request can be an IP address assigned to the sending interface, or an unspecified address if no address is assigned to the sending interface.

The routers in the network respond with an RA. The RAs can also be sent at periodic intervals. The RA contains the network part of the Layer 3 IPv6 address (IPv6 Prefix). The host uses the IPv6 prefix provided by the RA; it generates the universally unique host part of the address (interface identifier), and combines the two to derive the complete address. To establish continuous connectivity to the default router, the host starts the neighbor reachability state machine for the router.



ArubaOS uses Radvd, an open source Linux IPv6 Router Advertisement daemon maintained by Litech Systems Design.

You can perform the following tasks on the controller to enable, configure, and view the IPv6 RA status on a VLAN interface:

- Configure IPv6 RA on a VLAN
- Configure Optional Parameters for RA
  - Configure neighbor discovery reachable time
  - Configure neighbor discovery retransmit time
  - Configure RA DNS
  - Configure RA hop-limit
  - Configure RA interval
  - Configure RA lifetime
  - Configure RA managed configuration flag
  - Configure RA MTU
  - Configure RA other configuration flag
  - Configure RA Preference
  - Configure RA prefix
- View IPv6 RA Status

Configuring an IPv6 RA on a VLAN

You must configure the IPv6 RA functionality on a VLAN for it to send solicited/unsolicited router advertisements on the IPv6 network. You must configure the following for the IPv6 RA to be operational on a VLAN:

- IPv6 global unicast address
- enable IPv6 RA
- IPv6 RA prefix



- The advertised IPv6 prefix length must be 64 bits for the stateless address autoconfiguration to be operational.
- You can configure up to three IPv6 prefixes per VLAN interface.
- Each IPv6 prefix must have an on-link interface address configured on the VLAN.
- Ensure you configure the upstream routers to route the packets back to Dell controller.

You can use the WebUI or CLI to configure the IPv6 RA on a VLAN.

Using WebUI

1. Navigate to the **Configuration > Network > IP** page and select the **IP Interfaces** tab.
2. Edit a VLAN # and select **IP version** as *IPv6*.
3. To configure an IPv6 global unicast address, follow the steps below:
   a. Under **Details,** enter the IPv6 address and the prefix-length in the **IP Address/Prefix-length** field.
   b. (Optional) Select the **EUI64 Format** check box, if applicable.
   c. Click **Add** to add the address to the global address list.
4. To enable an IPv6 RA on a VLAN, select the **Enable Router Advertisements (RA)** check box under **Neighbor Discovery**.
5. To configure an IPv6 RA prefix for a VLAN, follow the steps below:
   a. Under **Neighbor Discovery**, enter an IPv6 prefix in the **IPv6 RA Prefix** field.
   b. Click **Add** to configure an IPv6 prefix for the VLAN.
   You can add up to three IPv6 prefixes per VLAN interface.
6. Click **Apply**.

Using CLI

Execute the following commands to configure router advertisements on a VLAN:

```
(host)(config) #interface vlan <vlanid>
(host)(config-subif)#ipv6 address <prefix>/<prefix-length>
(host)(config-subif)#ipv6 nd ra enable
(host)(config-subif)#ipv6 nd ra prefix X:X:X:X::X/64
```

Configuring Optional Parameters for RAs

In addition to enabling the RA functionality, you can configure the following IPv6 neighbor discovery and RA options on a VLAN:

- Neighbor discovery reachable time – the time, in milliseconds, that a node assumes a neighbor is reachable after receiving a reachability confirmation.
- Neighbor discovery retransmit time – the time, in milliseconds, between retransmitted Neighbor Solicitation messages.
- RA DNS – the IPv6 recursive DNS Server for the VLAN.



- On Linux systems, clients must run the open rdnssd daemon to support the DNS server option.
- Windows 7 does not support the DNS server option.

- RA hop-limit – the IPv6 RA hop-limit value. It is the default value to be placed in the Hop Count field of the IP header for outgoing (unicast) IP packets.
- RA interval – the maximum and minimum time interval between sending unsolicited multicast router advertisements from the interface, in seconds.
- RA lifetime – the lifetime associated with the default router in seconds. A value of zero indicates that the router is not a default router and will not appear on the default router list. The router lifetime applies only to the router's usefulness as a default router; it does not apply to information contained in other message fields or options.
- RA managed configuration flag (Enable DHCP for address) – a flag that indicates that the hosts can use the DHCP server for address autoconfiguration besides using RAs.

- RA maximum transmission unit (MTU) – the maximum transmission unit that all the nodes on a link use.
- RA other configuration flag (Enable DHCP for other information – a flag that indicates that the hosts can use the administered (stateful) protocol for autoconfiguration of other (non-address) information.
- RA preference – the preference associated with the default router.

You can use the WebUI or CLI to configure these options.



It is recommended that you retain the default value of the RA interval to achieve better performance.



If you enable RAs on more than 100 VLAN interfaces, some of the interfaces may not send out the RAs at regular intervals.

In the WebUI

1. Navigate to the **Configuration > Network > IP** page.
2. Select the **IP Interfaces** tab.
3. Edit the VLAN on which you want to configure the neighbor discovery or RA options.
4. Select **IP Version** as *IPv6*.
5. Under **Neighbor Discovery**, configure the following neighbor discovery and RA options for the VLAN based on your requirements:
   a. Enter a value in the **Reachable Time** field. The allowed range is 0-3,600,000 msec. The default value is zero.
   b. Enter a value in the **Retransmit Time** field. The allowed range is 0-3,600,000 msec.he default value is zero.
   c. Enter a DNS server name in the **IPv6 Recursive DNS Server** field.
   d. Enter a hop-limit value in the **RA hop-limit** field. The allowed range is 1-255. The default value is 64.
   e. Enter the maximum interval value in the **RA Interval(sec)** field. Allowed range is 4-1800 seconds. Default value is 600 seconds.
   f. Enter a value in the **RA Minimum Interval(sec)** field. Allowed range is 3-0.75 times the maximum RA interval value in seconds. The default minimum value is 0.33 times the maximum RA interval value
   g. Enter a value in the **RA Lifetime** field. A value of zero indicates that the router is not a default router. Apart from a zero value, the allowed range for the lifetime value is the RA interval time to 9,000 seconds. The default and minimum value is three times the RA interval time.
   h. Select the **DHCP for address** check box to enable the hosts to use the DHCP server for address autoconfiguration apart from any addresses auto configured using the RA.
   i. Enter a value in the **RA MTU Option** option. The allowed range is 1,280-maximum MTU allowed for the link.
   j. Select the **DHCP for Other Address** check box to enable the hosts to use the DHCP server for autoconfiguration of other (non-address) information.
   k. Select the router preference as **High**, **Medium**, or **Low**.
6. Click **Apply.**

In the CLI

Execute the following CLI commands to configure the neighbor discovery and RA options for a VLAN interface:

To configure neighbor discovery reachable time:

```
(host)(config) #interface vlan <vlan-id>
(host)(config-subif)#ipv6 nd reachable-time <value>
```

To configure neighbor discovery retransmit time:

```
(host)(config-subif)#ipv6 nd retransmit-time <value>
```

To configure IPv6 recursive DNS server:

```
(host)(config-subif)#ipv6 nd ra dns X:X:X:X::X
```

To configure RA hop-limit:

```
(host)(config-subif)#ipv6 nd ra hop-limit <value>
```

To configure RA interval:

```
(host) (config-subif)#ipv6 nd ra interval <value> <min-value>
```

To configure RA lifetime:

```
(host)(config-subif)#ipv6 nd ra life-time <value>
```

To enable hosts to use DHCP server for stateful address autoconfiguration:

```
(host)(config-subif)#ipv6 nd ra managed-config-flag
```

To configure maximum transmission unit for RA:

```
(host)(config-subif)#ipv6 nd ra mtu <value>
```

To enable hosts to use DHCP server for other non-address stateful autoconfiguration:

```
(host)(config-subif)#ipv6 nd ra other-config-flag
```

To specify a router preference:

```
(host)(config-subif)#ipv6 nd ra preference [High | Low | Medium]
```

To view the IPv6 RA status on the VLAN interfaces:

```
(host) #show ipv6 ra status
```

RADIUS Over IPv6

ArubaOS provides support for RADIUS authentication server over IPv6. You can configure an IPv6 host or specify an FQDN that can resolve to an IPv6 address for RADIUS authentication. The RADIUS server is in IPv4 mode by default. You must enable the RADIUS server in IPv6 mode to resolve the specified FQDN to IPv6 address.



You can only configure the global IPv6 address as the host for the Radius server in IPv6 mode.

You can configure the IPv6 host for the RADIUS server using the WebUI or CLI.

In the CLI

You must enable the `enable-ipv6` parameter to configure the RADIUS server in IPv6 mode.

```
(host)(config) #aaa authentication-server radius IPv6
(host)(RADIUS Server "IPv6") #enable-ipv6
```

Configure an IPv6 address as the host for RADIUS server using the following command:

```
(host)(RADIUS Server "IPv6") #host <ipv6-address>
```

The <host> parameter can also be a fully qualified domain name that can resolve to an IPv6 address.

NOTE

To resolve FQDN, you must configure the DNS server name using the `ip name-server <ip4addr>` command.

You can configure an IPv6 address for the NAS-IP parameter using the following CLI command:

```
(host) (RADIUS Server "Ipv6") #nas-ip6 <IPv6 address>
```

You can configure an IPv6 address for the Source Interface parameter using the following CLI command:

```
(host) (RADIUS Server "Ipv6") # source-interface vlan <vland-id> ip6addr <ip6addr>
```

Use the following CLI command to configure an IPv6 address for the global NAS IP which the controller uses to communicate with all the RADIUS servers:

```
(host) (config) #ipv6 radius nas-ip6 <IPv6 address>
```

You can also configure an IPv6 global source-interface for all the RADIUS server requests using the following commands:

```
(host)(config) #ipv6 radius source-interface loopback
(host)(config) #ipv6 radius source-interface vlan <vlan-id> <ip6addr>
```

In the WebUI

To configure an IPv6 host for a RADIUS server:

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **RADIUS Server** to display the RADIUS server List.
3. Select the required RADIUS server from the list to go to the Radius server page.
4. To enable the RADIUS server in IPv6 mode select the **Enable IPv6** check box.
5. To configure an IPv6 host for the selected RADIUS server specify an IPv6 address or an FQDN in the **Host** field.
6. Click **Apply**.

To configure an IPv6 address for the NAS-IP:

1. Select the **Advanced** tab.
2. Specify an IPv6 address in the **NAS IPv6** field.
3. Click **Apply**.

To configure an IPv6 global source-interface:

1. Select the **Advanced** tab.
2. To configure the IPv6 loopback interface as the source interface, select loopback from the **Source Interface v6** drop-down list.
3. To configure a VLAN interface as the source interface, specify the VLAN interface and the IPv6 address in the **Source Interface v6** field.
4. Click **Apply**.

TACACS Over IPv6

ArubaOS provides support for TACACS authentication server over IPv6. You can configure the global IPv6 address as the host for TACACS authentication using CLI or WebUI.

In the CLI

```
(host)(config) #aaa authentication-server tacacs IPv6
(host)(TACACS Server "IPv6") #host <ipv6-address>
```

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select**TACACS Server** to display the Server List.
3. Select the required server from the list to go to the TACACS server page.
4. To configure an IPv6 host for the selected server, specify an IPv6 address in the Host field.
5. Click **Apply**.

DHCPv6 Server

The DHCPv6 server enables network administrators to configure stateful/stateless options and manage dynamic IPv6 users connecting to a network. You can also configure domain name server using DHCPv6.

You can configure IPv6 pools with various configurations such as lease duration, DNS server, vendor specific options, and user defined options using DHCPv6. You can also exclude IPv6 addresses from subnets. Controller IPv6 addresses, VLAN interface IPv6 addresses, and DNS server addresses are excluded from use by default.

Similar to DHCPv4, a DHCPv6 server pool is associated with a VLAN only through the IPv6 address configured in that VLAN interface. A VLAN interface can have a maximum of three global unicast addresses, but only one DHCPv6 pool.

DHCPv6 server supports stateless configuration of clients with options apart from the network addresses described in RFC 3736.

Points to Remember

- Similar to IPv4, the default router configuration is not required for IPv6 pools as IPv6 compliant routers will send RAs. The RA source address will be the default-gateway for the clients.
- ArubaOS does not support DHCPv6 relay and Hospitality feature on DHCPv6.

DHCP Lease Limit

The following table provides the maximum number of DHCP leases (both v4 and v6) supported per controller platform:



There is a new enforcement to the existing DHCP limit during configuration.

Table 31: *DHCP Lease Limits*

Platform	Maximum number of DHCP Leases Supported
W-620	256
W-650	512
W-3200	512
W-6000M3	512
W-3400	512

Platform	Maximum number of DHCP Leases Supported
W-3600	512
W-7005	512
W-7010	1024
W-7030	2048
W-7210	5120
W-7220	10240
W-7240	15360

Configuring DHCPv6 Server

You must enable the global DHCPv6 knob for the DHCPv6 functionality to be operational. You can enable and configure DHCPv6 server using the WebUI or CLI.

In the WebUI

1. Navigate to **Configuration > Network > IP** page and select the **DHCP Server** tab.
2. Select the **IPv6 DHCP Server** check box to enable DHCPv6 globally.
3. If there are addresses that should not be assigned in the subnetwork:
   a. Under **Excluded Address Range**, click **Add** to create a list of IPv6 excluded address.
   b. Enter the excluded IPv6 address range in **IPv6 Excluded Range** and click **Done**. The specified address range gets added to the **IPv6 Excluded Address** list box. The starting IP address in the **Exclude Address Range** should always contain a unique value, if the IP address is already present, then the existing IP address is replaced with a new one, and a warning is displayed.
   c. Click **Apply**.
4. Under **Pool Configuration**, click **Add** to create a new DHCP server pool or click **Edit** to modify an existing DHCP server pool.



To enable the DHCPv6 Server functionality on an interface, select the **IP Interfaces** tab, edit the VLAN interface, and select a DHCP pool from the drop-down list under the **DHCP server** section. Ensure that the IP version of the VLAN interface is IPv6.

5. Select **IP Version** as **IPv6** to create a DHCPv6 pool.
6. Enter a name in **Pool Name** to configure an IPv6 pool name.
7. Enter an IPv6 address in **DNS Servers** to configure an IPv6 DNS server.



To configure multiple DNS servers, enter the IPv6 addresses separated by space.

8. Enter a value in **Domain Name** to configure the domain name.
9. Enter the number of days, hours, minutes, and seconds in **Lease** to configure the lease time. The default value is 12 hours.
10. Specify an IPv6 prefix in **Network** to configure an IPv6 network.
11. Enter the following details under **Option** to configure client specific DHCPv6 options.

a. Specify the option code in **Option**.
b. Select **IP** or **text** from the **IP/Text** drop-down list.
c. Enter a value in **Value**. If you selected *IP* in *step b*, then you must enter a valid IPv6 address in this field.
d. Click **Add**.

12. Click **Apply**.

In the CLI

To enable the DHCPv6 service you can use the following command:

```
(host)(config) #service dhcpv6
```

To configure a domain name server, execute the following commands:

```
(host)(config) #ipv6 dhcp pool <pool-name>
(host)(config-dhcpv6)#dns-server <ipv6-address>
```

To configure a domain name, use the following command:

```
(host)(config-dhcpv6)#domain-name <domain>
```

To configure DHCPv6 lease time, use the following command:

```
(host)(config-dhcpv6)#lease <days> <hours> <minutes> <seconds>
```

The default value is 12 hours.

To configure a DHCP network, use the following command:

```
(host)(config-dhcpv6)#network <network-prefix>
```

To configure a client specific option, use the following command:

```
(host)(config-dhcpv6)#option <code> [ip <ipv6-address> | text <string>]
```

To configure DHCP server preference, use the following command:

```
(host)(config-dhcpv6)#preference <value>
```

To enable DHCPv6 Server functionality on an interface, use the following command:

```
(host) (config) #interface vlan <vlan-id>
(host) (config-subif) #ipv6 dhcp server <pool-name>
```



The configured DHCPv6 pool subnet must match the interface prefix for DHCPv6 Server to be active.

To configure the IPv6 excluded address range for the DHCPv6 server, use the following command:

```
(host)(config)#ipv6 dhcp excluded-address <low-address> [<high-address>]
```

You can view the DHCPv6 server settings, statistics, and binding information using the CLI.

To view the DHCPv6 database, use the following command:

```
(host)#show ipv6 dhcp database
```

You can also view the DHCPv6 database for a specific pool, use the following command:

```
(host) (config) #show ipv6 dhcp database [pool <pool-name>]
(host) (config) #show ipv6 dhcp database pool DHCPv6
```

To view the DHCPv6 binding information, use the following command:

```
(host)# show ipv6 dhcp binding
```

To clear all the DHCPv6 bindings, use the following command:

```
(host)# clear ipv6 dhcp binding
```

To view the DHCPv6 server statistics, use the following command:

```
(host)(config) #show ip dhcp statistics
```

To view the DHCPv6 active pools, use the following command:

```
(host) #show ipv6 dhcp active-pools
```

Understanding ArubaOS Supported Network Configuration for IPv6 Clients

ArubaOS provides wired or wireless clients using IPv6 addresses with services such as firewall functionality, layer-2 authentication, and, with the installation of the Policy Enforcement Firewall Next Generation (PEFNG), identity-based security. The Dell controller does not provide routing or Network Address Translation to IPv6 clients (see Understanding IPv6 Exceptions and Best Practices on page 199).

Supported Network Configuration

Clients can be wired or wireless and use IPv4 and/or IPv6 addresses. An external IPv6 router is recommended for a complete routing experience (dynamic routing). You can use the WebUI or CLI to display IPv6 client information.

On the controller, you can configure both IPv4 and IPv6 client addresses on the same VLAN.

Understanding the Network Connection Sequence for Windows IPv6 Clients

This section describes the network connection sequence for Windows Vista/XP clients that use IPv6 addresses, and the actions performed by the AP and the controller.

1. The IPv6 client sends a Router Solicit message through the AP. The AP passes the Router Solicit message from the IPv6 client through the GRE tunnel to the controller.
2. The controller removes the 802.11 frame and creates an 802.3 frame for the Router Solicit message.
   a. The controller authenticates the user, applies firewall policies, and bridges the 802.3 frame to the IPv6 router.
   b. The controller creates entries in the user and session tables.
3. The IPv6 router responds with a Router Advertisement message.
4. The controller applies firewall policies, then creates an 802.11 frame for the Router Advertisement message. The controller sends the Router Advertisement through the GRE tunnel to the AP.
5. The IPv6 client sends a Neighbor Solicitation message.
6. The IPv6 router responds with a Neighbor Advertisement message.
7. If the DHCP is required to provide IPv6 addresses, the DHCPv6 process is started.
8. The IPv6 client sends data.
9. The controller removes the 802.11 frame and creates an 802.3 frame for the data.

   The controller authenticates the user, applies firewall policies and bridges the 802.3 frame to the IPv6 router. The controller creates entries in the user and session tables.



A client can have an IPv4 address and an IPv6 address, but the controller does not relate the states of the IPv4 and the IPv6 addresses on the same client. For example, if an IPv6 user session is active on a client, the controller will delete an IPv4 user session on the same client if the idle timeout for the IPv4 session is reached.

Understanding ArubaOS Authentication and Firewall Features that Support IPv6

This section describes ArubaOS features that support IPv6 clients.

Understanding Authentication

This release of ArubaOS only supports 802.1x authentication for IPv6 clients. You cannot configure layer-3 authentications to authenticate IPv6 clients.

Table 32: *IPv6 Client Authentication*

Authentication Method	Supported for IPv6 Clients?
802.1x	Yes
Stateful 802.1x (with non-Dell APs)	Yes
Local database	Yes
Captive Portal	Yes
VPN	No
xSec	No (not tested)
MAC-based	Yes

You configure 802.1x authentication for IPv6 clients in the same way as for IPv4 client configurations. For more information about configuring 802.1x authentication on the controller, see 802.1X Authentication on page 253.



This release does not support authentication of management users on IPv6 clients.

Working with Firewall Features

If you installed a Policy Enforcement Firewall Next Generation (PEFNG) license in the controller, you can configure firewall functions for IPv6 client traffic. While these firewall functions are identical to firewall functions for IPv4 clients, you need to explicitly configure them for IPv6 traffic. For more information about firewall policies, see Understanding Global Firewall Parameters on page 377.



Voice-related and NAT firewall functions are not supported for IPv6 traffic.

Table 33: *IPv6 Firewall Parameters*

Parameter	Description
Monitor Ping Attack (per 30 seconds)	Number of ICMP pings per 30 second, which if exceeded, can indicate a denial of service attack. Valid range is 1-16384 pings per 30 seconds. Recommended value is 120. Default: No default
Monitor TCP SYN Attack rate (per 30 seconds)	Number of TCP SYN messages per 30 second, which if exceeded, can indicate a denial of service attack. Valid range is 1-16384 pings per 30 seconds. Recommended value is 960. Default: No default
Monitor IP Session Attack (per 30 seconds)	Number of TCP or UDP connection requests per 30 second, which if exceeded, can indicate a denial of service attack. Valid range is 1-16384 requests per 30 seconds. Recommended value is 960.

Table 33: *IPv6 Firewall Parameters*

Parameter	Description
	Default: No default
Deny Inter User Bridging	Prevents the forwarding of Layer-2 traffic between wired or wireless users. You can configure user role policies that prevent Layer-3 traffic between users or networks but this does not block Layer-2 traffic. This option can be used to prevent traffic, such as Appletalk or IPX, from being forwarded. Default: Disabled
Deny All IP Fragments	Drops all IP fragments. **NOTE:** Do not enable this option unless instructed to do so by a Dell representative. Default: Disabled
Enforce TCP Handshake Before Allowing Data	Prevents data from passing between two clients until the three-way TCP handshake has been performed. This option should be disabled when you have mobile clients on the network, as enabling this option will cause mobility to fail. You can enable this option if there are no mobile clients on the network. Default: Disabled
Prohibit IP Spoofing	Enables detection of IP spoofing (where an intruder sends messages using the IP address of a trusted client). When you enable this option, IP and MAC addresses are checked for each ARP request/response. Traffic from a second MAC address using a specific IP address is denied, and the entry is not added to the user table. Possible IP spoofing attacks are logged and an SNMP trap is sent. Default: Disabled
Prohibit RST Replay Attack	When enabled, closes a TCP connection in both directions if a TCP RST is received from either direction. You should not enable this option unless instructed to do so by a Dell representative. Default: Disabled
Session Mirror Destination	Destination (IPv4 address or controller port) to which mirrored session packets are sent. You can configure IPv6 flows to be mirrored with the session ACL "mirror" option. This option is used only for troubleshooting or debugging. Default: N/A
Session Idle Timeout	Set the time, in seconds, that a non-TCP session can be idle before it is removed from the session table. Specify a value in the range 16–259 seconds. You should not set this option unless instructed to do so by a Dell representative. Default: 30 seconds
Per-packet Logging	Enables logging of every packet if logging is enabled for the corresponding session rule. Normally, one event is logged per session. If you enable this option, each packet in the session is logged. You should not enable this option unless instructed to do so by a Dell representative, as doing so may create unnecessary overhead on the controller. Default: Disabled (per-session logging is performed)
IPv6 Enable	Enables IPv6 globally.

The following examples configure attack rates and the session timeout for IPv6 traffic.

To configure the firewall function via the WebUI:

1. Navigate to the **Configuration > Advanced Services > Stateful Firewall > Global Setting** page.
2. Under the **IPv6** column, enter the following:
   - For **Monitor Ping Attack**, enter **15**
   - For **Monitor IP Session Attack**, enter **25**

- For **Session Idle Timeout**, enter **60**

3. Click **Apply**.

To configure firewall functions using the command line interface, issue the following commands in config mode:

```
ipv6 firewall attack-rate ping 15
ipv6 firewall attack-rate session 25
ipv6 firewall session-idle-timeout 60
```

Understanding Firewall Policies

A user role, which determines a client's network privileges, is defined by one or more firewall policies. A firewall policy consists of rules that define the source, destination, and service type for specific traffic, and whether you want the controller to permit or deny traffic that matches the rule.

You can configure firewall policies for IPv4 traffic or IPv6 traffic, and apply IPv4 and IPv6 firewall policies to the same user role. For example, if you have employees that use both IPv4 and IPv6 clients, you can configure both IPv4 and IPv6 firewall policies and apply them both to the "employee" user role.

The procedure to configure an IPv6 firewall policy rule is similar to configuring a firewall policy rule for IPv4 traffic, but with some differences. Table 18 describes the required and optional parameters for an IPv6 firewall policy rule.

Table 34: *IPv6 Firewall Policy Rule Parameters*

Field	Description
Source (required)	Source of the traffic: • **any**: Acts as a wildcard and applies to any source address. • **user**: This refers to traffic from the wireless client. • **host**: This refers to traffic from a specific host. When this option is chosen, you must configure the IPv6 address of the host. For example, 2002:d81f:f9f0:1000:c7e:5d61:585c:3ab. • **network**: This refers to a traffic that has a source IP from a subnet of IP addresses. When you chose this option, you must configure the IPv6 address and network mask of the subnet. For example, 2002:ac10:fe:: ffff:ffff:ffff::. • **alias**: This refers to using an alias for a host or network. **NOTE:** This release does not support IPv6 aliases. You cannot configure an alias for an IPv6 host or network.
Destination (required)	Destination of the traffic, which you can configure in the same manner as Source.
Service (required)	**NOTE:** Voice over IP services are unavailable for IPv6 policies. Type of traffic: • **any**: This option specifies that this rule applies to any type of traffic. • **tcp**: Using this option, you configure a range of TCP port(s) to match the rule to be applied. • **udp**: Using this option, you configure a range of UDP port(s) to match the rule to be applied. • **service**: Using this option, you use one of the pre-defined services (common protocols such as HTTPS, HTTP, and others) as the protocol to match the rule to be applied. You can also specify a network service that you configure by navigating to the **Configuration > Advanced Services > Stateful Firewall > Network Services** page. • **protocol**: Using this option, you specify a different layer 4 protocol (other than TCP/UDP) by configuring the IP protocol value.
Action (required)	The action that you want the controller to perform on a packet that matches the specified criteria.

Table 34: *IPv6 Firewall Policy Rule Parameters*

Field	Description
	• **permit:** Permits traffic matching this rule. • **drop:** Drops packets matching this rule without any notification. **NOTE:** The only actions for IPv6 policy rules are permit or deny; in this release, the controller cannot perform network address translation (NAT) or redirection on IPv6 packets. You can specify options such as logging, mirroring, or blacklisting (described below).
Log (optional)	Logs a match to this rule. This is recommended when a rule indicates a security breach, such as a data packet on a policy that is meant only to be used for voice calls.
Mirror (optional)	Mirrors session packets to a datapath or remote destination specified in the IPv6 firewall function (see "Session Mirror Destination" in Table 33). If the destination is an IP address, it must be an IPv4 IP address.
Queue (optional)	The queue in which a packet matching this rule should be placed. Select **High** for higher priority data, such as voice, and **Low** for lower priority traffic.
Time Range (optional)	Time range for which this rule is applicable. You configure time ranges in the **Configuration > Security > Access Control > Time Ranges** page.
Black List (optional)	Automatically blacklists a client that is the source or destination of traffic matching this rule. This option is recommended for rules that indicate a security breach where the blacklisting option can be used to prevent access to clients that are attempting to breach the security.
TOS (optional)	Value of type of service (TOS) bits to be marked in the IP header of a packet matching this rule when it leaves the controller.
802.1p Priority (optional)	Value of 802.1p priority bits to be marked in the frame of a packet matching this rule when it leaves the controller.

The following example creates a policy "ipv6-web-only" that allows only web (HTTP and HTTPS) access for IPv6 clients and assigns the policy to the user role "web-guest."



The user role "web-guest" can include both IPv6 and IPv4 policies, although this example only shows configuration of an IPv6 policy.

Creating an IPv6 Firewall Policy

Following the procedure below to create an IPv6 firewall policy via the WebUI.

1. Navigate to the **Configuration > Security > Access Control > Policies** page.
2. Click **Add** to create a new policy.
3. Enter **ipv6-web-only** for the Policy Name.
4. To configure a firewall policy, select **Session** for Policy Type.
5. Click **Add** to add a rule that allows HTTP traffic.
   a. Under **IP Version** column, select **IPv6**.
   b. Under **Source**, select **network** from the drop-down list.
   c. For **Host IP**, enter **2002:d81f:f9f0:1000::**.
   d. For **Mask**, enter **64** as the prefix-length.
   e. Under **Service**, select **service** from the drop-down list.
   f. Select **svc-http** from the scrolling list.

g. Click **Add**.

6. Click **Add** to add a rule that allows HTTPS traffic.
   a. Under **IP Version** column, select **IPv6.**
   b. Under **Source,** select **network** from the drop-down list.
   c. For **Host IP**, enter **2002:d81f:f9f0:1000::**.
   d. For **Mask**, enter **64** as the prefix-length.
   e. Under **Service**, select **service** from the drop-down list.
   f. Select **svc-https** from the scrolling list.
   g. Click **Add**.



Rules can be reordered using the up and down arrow buttons provided for each rule.

7. Click **Apply**. The policy is not created until the configuration is applied.

To create an IPv6 firewall policy using the command-line interface, issue the following commands in config mode:

```
ip access-list session ipv6-web-only
  ipv6 network 2002:d81f:f9f0:1000::/64 any svc-http permit
  ipv6 network 2002:d81f:f9f0:1000::/64 any svc-https permit
```

Assigning an IPv6 Policy to a User Role

To assign an IPv6 policy using the WebUI:

1. Navigate to the **Configuration > Security > Access Control > User Roles** page.
2. Click **Add** to create a new user role.
3. Enter **web-guest** for Role Name.
4. Under Firewall Policies, click **Add**. From Choose from Configured Policies, select the “ipv6-web-only” IPv6 session policy from the list.
5. Click **Done** to add the policy to the user role.
6. Click **Apply**.

To assign an IPv6 policy to a user role via the command-line interface, issue the following command in config mode:

```
user-role web-guest
  access-list session ipv6-web-only position 1
```

Understanding DHCPv6 Passthrough/Relay

The controller forwards DHCPv6 requests from IPv6 clients to the external IPv6 router. On the external IPv6 router, you must configure the controller’s IP address as the DHCP relay. You do *not* need to configure an IP helper address on the controller to forward DHCPv6 requests.

Managing IPv6 User Addresses

Viewing or Deleting User Entries

To view or delete IPv6 user entries via the WebUI:

1. Navigate to the **Monitoring > Controller > Clients** page.

2. Click the **IPv6** tab to display IPv6 clients.
3. To delete an entry in the IPv6 client display, click the radio button to the left of the client and then click **Disconnect**.

To view user entries for IPv6 clients using the command line interface, use the **`show user-table`** command in enable mode. To delete a user entry for an IPv6 client, access the CLI in config mode and use the **`aaa ipv6 user delete`** command. For example:

```
aaa ipv6 user delete 2002:d81f:f9f0:1000:e409:9331:1d27:ef44
```

Understanding User Roles

An IPv6 user or a client can inherit the corresponding IPv4 roles. A user or client entry on the user table will contain the user or client's IPv4 and IPv6 entries. After captive-portal authentication, a IPv4 client can acquire a different role. This role is also updated on the client's IPv6 entry in the user table.

Viewing Datapath Statistics for IPv6 Sessions

To view datapath session statistics for individual IPv6 sessions, access the command-line interface in enable mode and issue the command **`show datapath session ipv6`**. To display the user entries in the datapath, access the command-line interface in enable mode, and issue the command **`show datapath user ipv6`**. For details on each of these commands and the output they display, refer to the *Dell Networking W-Series ArubaOS Command Line Reference Guide*.

Understanding IPv6 Exceptions and Best Practices

The IPv6 best practices are provided below:

- Ensure that you enable IPv6 globally.
- The uplink port must be trusted. This is the same behavior as IPv4.
- Ensure that the `validuser` session ACL does not block IPv6 traffic.
- There must not be any ACLs that drop ICMPv6 or DHCPv6 traffic. It is acceptable to drop DHCPv6 traffic if the deployment uses Stateless Address Auto Configuration (SLAAC) only.
- If an external device provides RA:
  - It is not recommended to advertise too many prefixes in RA.
  - The controller supports a maximum of four IPv6 user entries in the user table. If a client uses more than four IPv6 addresses at a time, the user table is refreshed with the latest four active entries without disrupting the traffic flow. However, this may have some performance impact.
- Enable **BCMC Optimization** under interface VLAN to drop any random IPv6 multicast traffic. DHCPv6, ND, NS, and RA traffic are not dropped when you enable this option.



It is recommended to enable **BCMC Optimization** only if mDNS traffic is not used in the network, as mDNS traffic gets dropped if this option is enabled.

- It is not recommended to enable preemption on the master redundancy model. If preemption is disabled and if there is a failover, the new primary controller remains the primary controller even when the original master is online again. The new primary controller does not revert to it's original state unless forced by the administrator. Disabling preemption prevents the master from "flapping" between two controllers and allows the administrator to investigate the cause of the outage.
- While selecting a source address, the number of common bits between each source address in the list, is checked from the left most bit. This is followed by selection of the source address that has the maximum number of matching bits with the destination address. If more than one source addresses has the same

number of matching bits with the destination address, the kernel selects that source address that is most recently configured on the system. It is essential that the administrator/user configures the network appropriately, if a particular VLAN interface needs to be selected as the source. For example, in case of Dot1x authentication the administrator/user can configure the source interface appropriately so that it is selected for authentication process. For more information on IPv6 source address selection, see **RFC 3848**.

ArubaOS does not support the following functions for IPv6 clients:

- The controller offers limited routing services to IPv6 clients, so it is recommended to use an external IPv6 router for a complete routing experience (dynamic routing).
- Vo IP ALG is not supported for IPv6 clients.
- Remote AP supports IPv6 clients in tunnel forwarding mode only. The Remote AP bridge and split-tunnel forwarding modes do not support IPv6 clients. Secure Thin Remote Access Point (STRAP) cannot support IPv6 clients.
- IPSec is not supported over IPv6.
- IPv6 Auto configuration and IPv6 Neighbor Discovery mechanisms does not apply to IPv6 tunnels.
- Tunnel Encapsulation Limit, Tunnel-group, and MTU discovery options on IPv6 tunnels are not supported.
- IPSec is not supported in this release, so IPv6 GRE cannot be used for master-local setup.

The ArubaOS implementation of Link Aggregation Control Protocol (LACP) is based on the standards specified in 802.3ad. LACP provides a standardized means for exchanging information, with partner systems, to form a link aggregation group (LAG). LACP avoids port channel misconfiguration.

Two devices (actor and partner) exchange LACP data units (DUs) when forming a LAG. Once multiple ports in the system have the same actor system ID, actor key, partner system ID, and partner key, they belong to the same LAG.

The maximum number of supported port-channels is eight. With the introduction of LACP, this number remains the same. A port-channel group (LAG) is created either statically or dynamically via LACP. This chapter contains the following topics:

- Understanding LACP Best Practices and Exceptions on page 201
- Configuring LACP on page 201
- LACP Sample Configuration on page 203



For information on configuring LACP on W-AP220 Series and W-AP270 Series access points, see Link Aggregation Support on W-AP220 Series and W-AP270 Series on page 529

Understanding LACP Best Practices and Exceptions

- LACP is disabled by default.
- LACP depends on periodical Tx/Rx of LACP data units (LACPDU). Any failures are noticed immediately and that port is removed from the LAG.
- The maximum LAG supported per system is 8 groups; each group can be created statically or via LACP.
- Each LAG can have up to 8 member ports.
- The LAG group identification (ID) range is 0–7 for both static (port-channel) and LACP groups.
- When a port is added to a LACP LAG, it inherits the port-channel's properties (in other words, VLAN membership, trunk status , and so on.)
- When a port is added to LACP LAG, the port's property (i.e. speed) is compared to the existing port properties. If there is a mismatch, the command is rejected.
- The LACP commands can not be configured on a port that is already a member of a static port-channel. Similarly, if the group assigned in the command **`lacp group`** `<number>` already contains static port members, the command is rejected.
- The port uses the group number as its actor admin key.
- All ports use long timeout values (90 seconds) by default.
- The output of the command **`show interface port-channel`** now indicates if the LAG is created by LACP (dynamic) or static configuration. If the LAG is created via LACP, you can not add/delete any ports under that port channel. All other commands are allowed.

Configuring LACP

Two LACP configured devices exchange LACPDUs to form a link aggregation group (LAG). A device is configurable as an active or passive participant. In active mode, the device initiates DUs irrespective of the

partner state; passive mode devices respond only to the incoming DUs sent by the partner device. Hence, to form a LAG group between two devices, one device must be an active participant. For detailed information on the LACP commands, see the Command-Line Interface *Reference Guide*.

In the CLI

LACPDUs exchange their corresponding system identifier/priority along with their port's key/priority. This information determines the LAG of a given port. The LAG for a port is selected based on its keys; the port is placed in that LAG only when its system ID/key and partner's system ID/key matches the other ports in the LAG (if the group has ports).

1. Enable LACP and configure the per-port specific LACP. The group number range is 0–7.

   lacp group <group_number> mode {active | passive}

   - Active mode—the interface is in an active negotiating state. LACP runs on any link that is configured to be in the active state. The port in an active mode also automatically initiates negotiations with other ports by initiating LACP packets.
   - Passive mode—the interface is *not* in an active negotiating state. LACP runs on any link that is configured in a passive mode. The port in a passive mode responds to negotiations requests from other ports that are in an active mode. Ports in passive mode respond to LACP packets.



A port in a passive mode cannot set up a port channel (LAG group) with another port in a passive mode.

2. Set the timeout for the LACP session. The timeout value is the amount of time that a port-channel interface waits for a LACPDU from the remote system before terminating the LACP session. The default long time-out value is 90 seconds; short is 3 seconds.

   ```
   lacp timeout {long | short}
   ```

3. Set the port priority.

   ```
   lacp port-priority <priority_value>
   ```

   The higher the priority value the lower the priority. The range is 1-65535 and the default is 255.

4. View your LACP configuration.

   The port uses the group number +1 as the "actor admin key". All the ports use the long timeout value (90 seconds) by default.

   ```
   (host)#show lacp 0 neighbor
   Flags:  S - Device is requesting Slow LACPDUs
               F - Device is requesting fast LACPDUs
               A - Device is in active mode P - Device is in passive mode
   Partner's information
   ---------------------
   Port     Flags  Pri  OperKey  State Num  Dev Id
   ----     -----  ---- -------  ----- ---- ----------------
   FE 1/1   SA     1    0x10     0x45  0x5  00:0b:86:51:1e:70
   FE 1/2   SA     1    0x10     0x45  0x6  00:0b:86:51:1e:70
   ```

   When a port in a LAG is misconfigured (the partner device is different than the other ports), or the neighbor timesout or can not exchange LACPDUs with the partner, the port status is displayed as "DOWN" (see the following example):

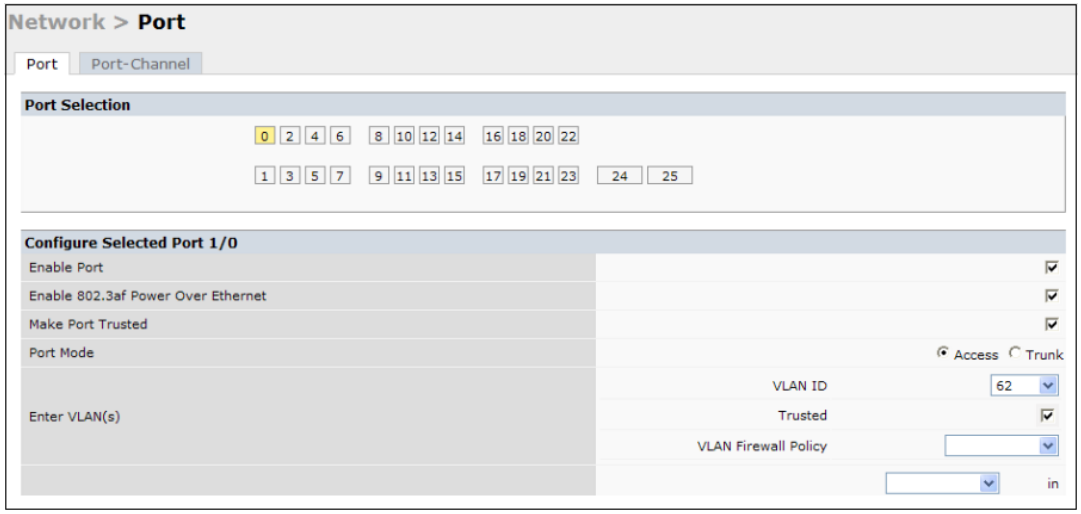
   ```
   (host)#show lacp 0 internal
   Flags: S - Device is requesting Slow LACPDUs
                           F - Device is requesting fast LACPDUs
                           A - Device is in active mode P - Device is in passive mode

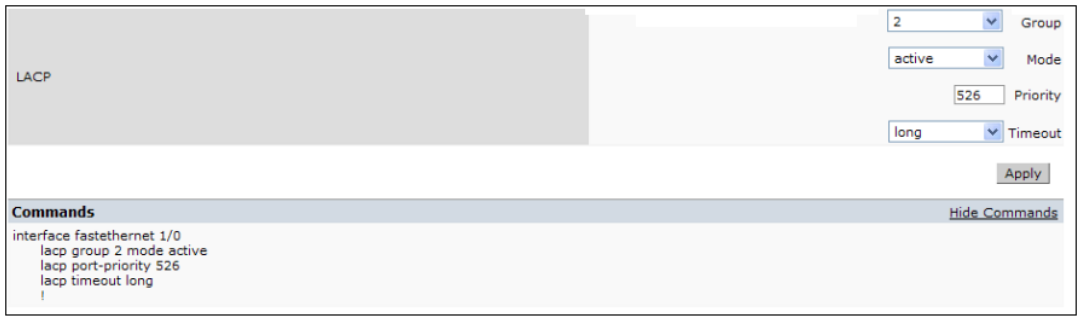
   Port     Flags  Pri  AdminKey  OperKey  State Num  Status
   ```

```
----    -----   ----  --------  --------  -----  ----  -------
FE 1/1  SA      1     0x1       0x1       0x45   0x2   DOWN
FE 1/2  SA      1     0x1       0x1       0x45   0x3   UP
```

In the WebUI

Access LACP from the **Configuration >Network >Port** tabs. Use the drop-down list to enter the LACP values.





- LACP Group— the link aggregation group (LAG) number; the range is 0 to 7.
- Mode— active negotiation state or not in an active negotiation state indicated by the *passive* option.
- Priority—the port priority value; the range is 1-65535 and the default is 255.
- Timeout— time out value for the LACP session. The long default is 90 seconds; the short default is 3 seconds.



For information on configuring LACP on W-AP220 Series and W-AP270 Series access points, see Link Aggregation Support on W-AP220 Series and W-AP270 Series on page 529

LACP Sample Configuration

The following sample configuration is for FastEthernet (FE) port/slot 1/0, 1/1, and 1/2:

```
interface fastethernet 1/0
        description "FE1/0"
        trusted vlan 1-4094
        lacp group 0 mode active
!
interface fastethernet 1/1
        description "FE1/1"
```

```
        trusted vlan 1-4094
        lacp timeout short
        lacp group 0 mode active
!
interface fastethernet 1/2
        description "FE1/2"
        trusted vlan 1-4094
        lacp group 0 mode passive
!
```

OSPFv2 (Open Shortest Path First) is a dynamic Interior Gateway routing Protocol (IGP) based on IETF RFC 2328. The OSPF uses the shortest or fastest routing path. Dell's implementation of OSPFv2 allows Dell controllers to deploy effectively in a Layer 3 topology. Dell controllers can act as default gateway for all clients and forward user packets to the upstream router. An Dell controller can be used for Instant AP VPN termination from the branch office, and the OSPF on the controller can be used to redistribute branch routes into corporate OSPF domain. The information on this chapter is in the following sections:

- Understanding OSPF Deployment Best Practices and Exceptions on page 205
- Understanding OSPFv2 by Example using a WLAN Scenario on page 206
- Understanding OSPFv2 by Example using a Branch Office Scenario on page 207
- Configuring OSPF on page 208
- Sample Topology and Configuration on page 210

Understanding OSPF Deployment Best Practices and Exceptions

OSPF is a robust routing protocol addressing various link types and deployment scenarios. The Dell implementation applies to two main use cases; WLAN Scenarios and Branch Office Scenario.

- OSPF is disabled by default.
- Dell controllers support only one OSPF instance.
- Convergence takes between 5 and 15 seconds.
- All area types are supported.
- Multiple configured areas are supported.
- A Dell controller can act as an ABR (Area border router).
- OSPF supports VLAN and GRE tunnel interfaces.
- To run OSPF over IPSec tunnels, a Layer 3 GRE tunnel is configured between two routers with GRE destination addresses as the inner address of the IPsec tunnel. OSPF is enabled on the Layer 3 GRE tunnel interface, and all of the OSPF control packets undergo GRE encapsulation before entering the IPsec tunnels.

  The default MTU value for a Layer 3 GRE tunnel in a Dell controller is 1100. When running OSPF over a GRE tunnel between a Dell controller and another vendor's router, the MTU values must be the same on both sides of the GRE tunnel.

  The following table provides information on the maximum OSPF routes supported for various platforms:

Table 35: *Maximum OSPF Routes*

Platform	Branches	Routes
W-3600	8K	8K
W-6000M3	8K	8K
W-7210	8K	8K

Platform	Branches	Routes
W-7220	16K	16K
W-7240	32K	32K

Below are some guidelines regarding deployment and topology for this release of OSPFv2.

- In the WLAN scenario, configure the Dell controller and all upstream routers in totally stub area; in the Branch Office scenario, configure as stub area so that the Branch Office controller can receive corporate subnets.
- In the WLAN scenario upstream router, only configure the interface connected to the controller in the same area as the controller. This will minimize the number of local subnet addresses advertised by the upstream router to the controller.
- Use the upstream router as the designated router (DR) for the link/interface between the controller and the upstream router.
- The default MTU value for a Layer 3 GRE tunnel in a Dell controller is 1100. When running OSPF over a GRE tunnel between a Dell controller and another vendor's router, the MTU values must be the same on both sides of the GRE tunnel.
- Do not enable OSPF on any uplink/WAN interfaces on the Branch Office Controller. Enable OSPF only on the Layer 3 GRE tunnel connecting the master controller.
- Use only one physical port in the uplink VLAN interface that is connecting to the upstream router. This will prevent broadcasting the protocol PDUs to other ports and hence limit the number of adjacencies on the uplink interface to only one.

Understanding OSPFv2 by Example using a WLAN Scenario

In the WLAN scenario, the Dell controller acts as a default gateway for all the clients, and talks to one or two upstream routers for redundancy. The controller advertises all the user subnet addresses as stub addresses to the routers via LSAs.



Totally stub areas *see* only default route and to the areas themselves.

WLAN Topology

The controller is configured with VLAN 10 and VLAN 12 as user VLANs. These VLANs have clients on the subnets, and the controller is the default router for those clients. VLAN 4 and VLAN 5 both have OSPF enabled. These interfaces are connected to upstream routers (Router 1 and Router 2). The OSPF interface cost on VLAN 4 is configured lower than VLAN 5. The IDs are:

- Dell controller— 40.1.1.1
- Router 1— 50.1.1.1
- Router 2— 60.1.1.1

Based on the cost of the uplink interface, the default route from one of the upstream routers is installed in the forwarding information base (FIB) by the routing information base/route table manager (RIB/RTM) module.

WLAN Routing Table

View the controller routing table using the **`show ip route`** command:

```
(host)#show ip route
```

```
Codes: C - connected, O - OSPF, R - RIP, S - static
       M - mgmt, U - route usable, * - candidate default
```

Below is the routing table for Router 1:

```
(router1) #show ip route

O    10.1.1.0/24  [1/0] via 4.1.1.1
O    12.1.1.0/24  [1/0] via 4.1.1.1
C    4.1.1.0 is directly connected, VLAN4
```

Below is the routing table for Router 2:

```
(router2) #show ip route

O    10.1.1.0/24  [2/0] via 5.1.1.1
O    12.1.1.0/24  [2/0] via 5.1.1.1
C    5.1.1.0 is directly connected, VLAN5
```

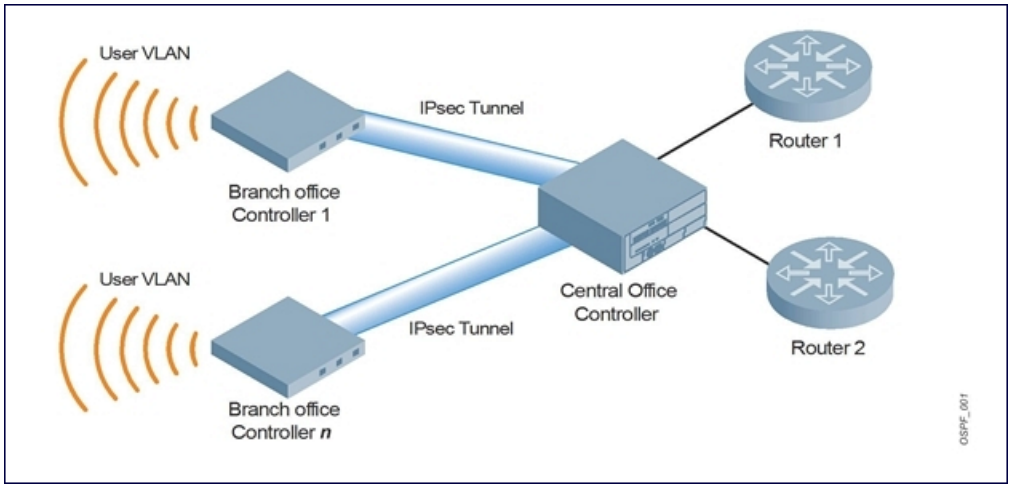
Understanding OSPFv2 by Example using a Branch Office Scenario

The branch office scenario has a number of remote branch offices with controllers talking to a central office via a concentrator/controller using site-to-site VPN tunnels or master-local IPsec tunnels. The central office controller is in turn talking to upstream routers (see Figure 22). In this scenario, the default route is normally pointed to the uplink router, in many cases the ISP. Configure the area as stub so that inter-area routes are also advertised enabling the branch office controller to reach the corporate subnets.

Branch Office Topology

All the OSPF control packets exchanged between the Branch office and the central office controllers undergo GRE encapsulation before entering the IPsec tunnels. The controllers in the branch offices advertise all the user subnet addresses to the Central office controller as stub addresses in router LSA. The central office controller in turn forwards those router LSAs to the upstream routers.

Figure 22 *Branch Office OSPF Topology*



All the branch office controllers, the Central office controller, and the upstream routers are part of a stub area. Because the OSPF packets follow GRE encapsulation over IPsec tunnels, the Central office controller can be a controller or any vendor's VPN concentrator. Regardless, the controller in the branch office will operate with other vendors seamlessly.

In Figure 22, the branch office controller is configured using VLAN 14 and VLAN 15. Layer 3 GRE tunnel is configured with IP address 20.1.1.1/24 and OSPF is enabled on the tunnel interface.

In the Central office controller, OSPF is enabled on VLAN interfaces 4, 5, and the Layer 3 GRE tunnel interface (configured with IP address 20.1.1.2/24). OSPF interface cost on VLAN 4 is configured lower than VLAN 5.

Branch Office Routing Table

View the branch office controller routing table using the **`show ip route`** command:

```
(host)#show ip route

Codes: C - connected, O - OSPF, R - RIP, S - static
       M - mgmt, U - route usable, * - candidate default
```

The routing table for the central office controller is below:

```
(host)#show ip route

Gateway of last resort is 4.1.1.2 to network 0.0.0.0

O*    0.0.0.0/0  [1/0] via 4.1.1.2*
O    14.1.1.0/24  [1/0] via 30.1.1.1*
O    15.1.1.0/24  [1/0] via 30.1.1.1*
C    4.1.1.0 is directly connected, VLAN4
C    5.1.1.0 is directly connected, VLAN5
C    20.1.1.0 is directly connected, Tunnel 1
```

The routing table for Router 1 is below:

```
(router1) #show ip route

O    14.1.1.0/24  [1/0] via 4.1.1.1
O    15.1.1.0/24  [1/0] via 4.1.1.1
C    4.1.1.0 is directly connected, VLAN4
```

The routing table for Router 2 is below:

```
(router2) #show ip route

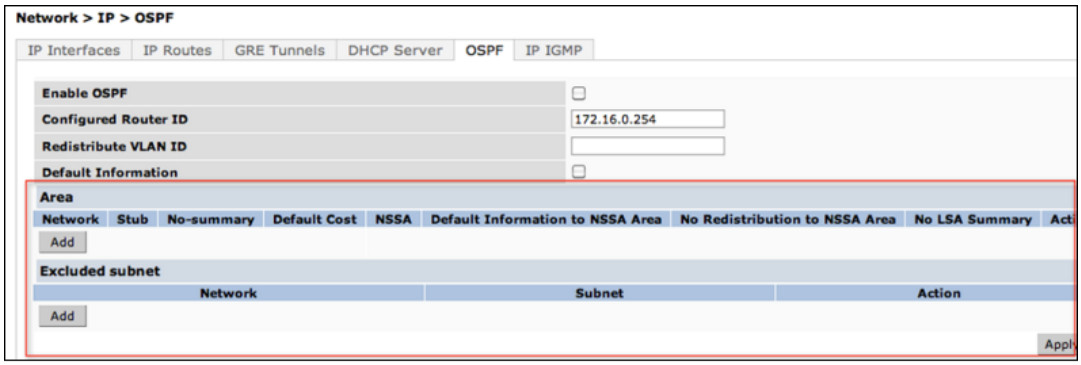
O    14.1.1.0/24  [1/0] via 5.1.1.1
O    15.1.1.0/24  [1/0] via 5.1.1.1
C    5.1.1.0 is directly connected, VLAN5
```

Configuring OSPF

To configure general OSPF settings from the OSPF tab, perform the following steps:

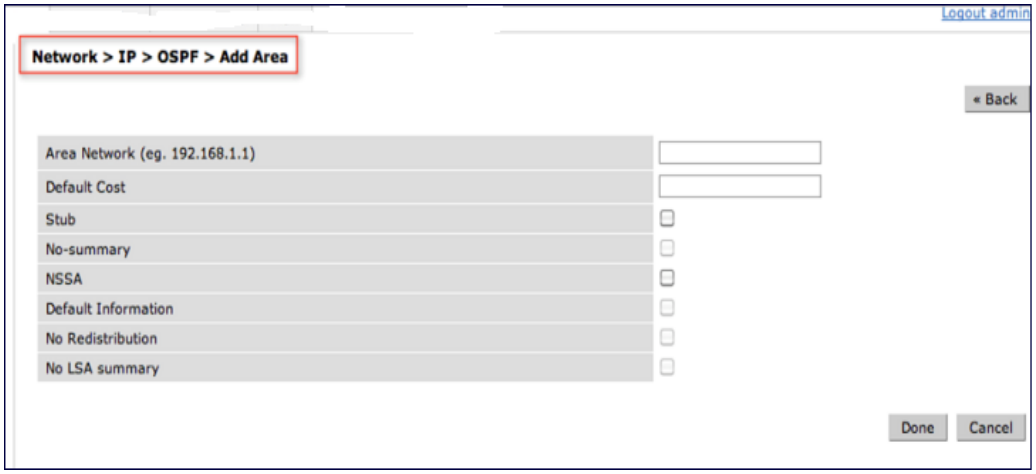
1. Navigate to the **Configuration >IP** page (see Figure 23). The Area and Excluded subnets are displayed in table format. If not explicitly specified for OSPF, the router ID defaults to the switch IP.

Figure 23 *General OSPF Configuration*



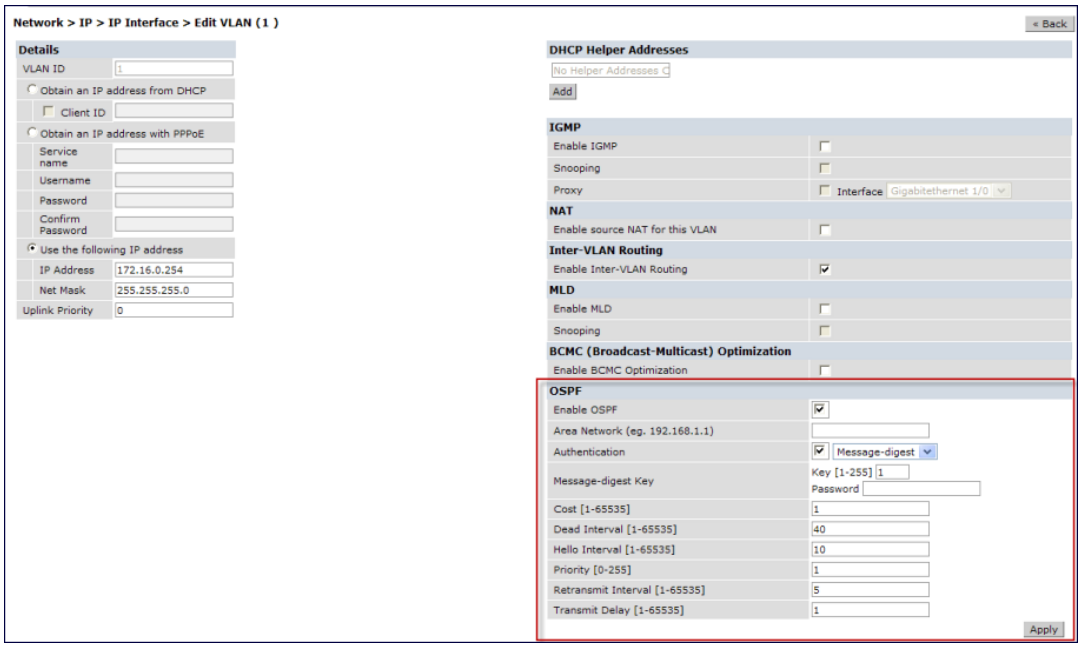
2. Click **Add** to add an area (see Figure 24).

Figure 24 *Add an OSPF Area*



3. Configure the OSPF interface settings in the Configuration screen (Figure 25). If OSPF is enabled, the parameters contain the correct default values. You can edit the OSPF values only when you enable OSPF on the interface.

Figure 25 *Edit OSPF VLAN Settings*



OSPF monitoring is available from an IP Routing sub-section (**Controller > IP Routing > Routing**). Both Static and OSPF routes are available in table format.

OSPF Interfaces and Neighboring information is available from the **OSPF** tab. The Interface information includes transmit (TX) and receive (RX) statistics.

Exporting VPN Client Addresses to OSPF

You can configure VPN client addresses so that they can be exported to OSPF and be advertised as host routes (/32). Exporting applies to any VPN client address regardless of how it is assigned.

In the WebUI

1. Navigate to the **Configuration > Advanced Services > All Profiles > VPN Authentication > default** page.
2. (Optional) Regardless of how an authentication server is contacted, the **Export VPN IP address as a route** option causes any VPN client address to be exported to OSPF using IPC. Note that the Framed-IP-Address attribute is assigned the IP address as long as any server returns the attribute. The Framed-IP-Address value always has a higher priority than the local address pool.
3. Click **Apply**.

In the CLI

Use this command to export the VPN client's assigned address to OSPF using IPC:

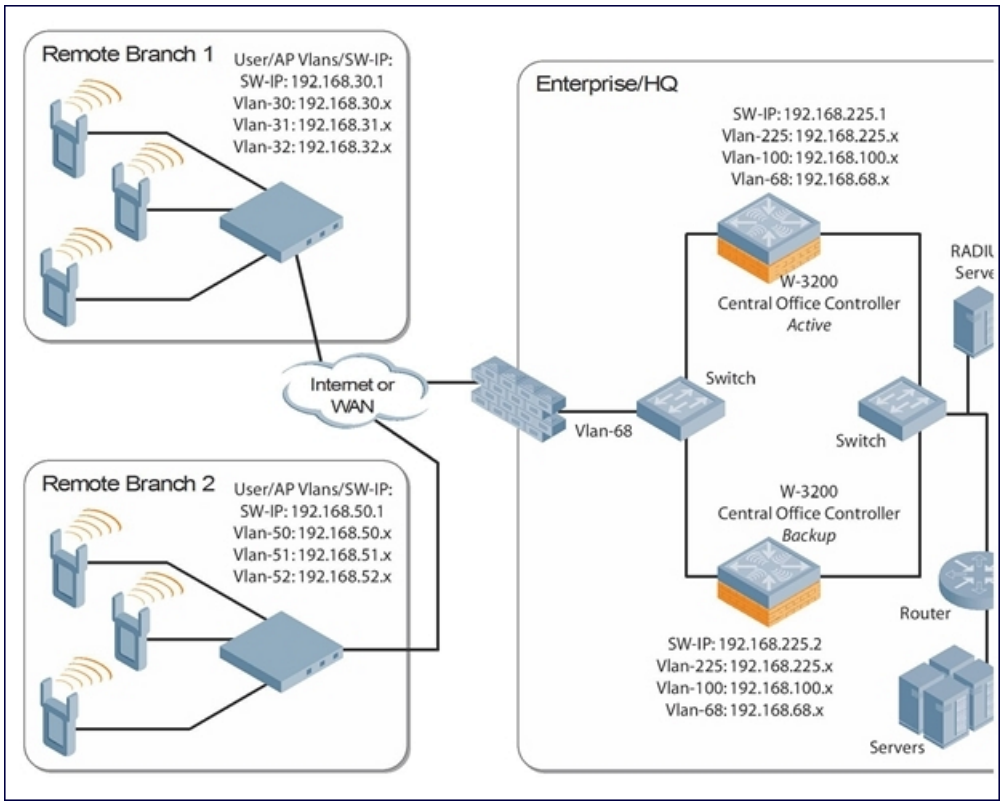
```
(host) (config) #aaa authentication vpn default
(host) (VPN Authentication Profile "default") #
(host) (VPN Authentication Profile "default") # export-route
```

Use the `show ip ospf` database command to show LSA types that are generated.

Sample Topology and Configuration

The figure below displays a sample OSPF topology followed by sample configurations of the Remote Branch 1, Remote Branch 2, and the W-3200 Central Office Controller (Active and Backup).

Figure 26 *Sample OSPF Topology*



Remote Branch 1

```
controller-ip vlan 30
vlan 16
vlan 30
vlan 31
vlan 32
interface gigabitethernet  1/0
        description "GE1/0"
        trusted
        switchport access vlan 16
!
interface gigabitethernet  1/1
        description "GE1/1"
        trusted
        switchport access vlan 30
!

interface gigabitethernet  1/2
        description "GE1/2"
        trusted
        switchport access vlan 31
!
interface gigabitethernet  1/3
        description "GE1/3"
        trusted
        switchport access vlan 32
!
interface vlan 16
        ip address 192.168.16.251 255.255.255.0
```

```
!
interface vlan 30
        ip address 192.168.30.1 255.255.255.0
!
interface vlan 31
        ip address 192.168.31.1 255.255.255.0
!
interface vlan 32
        ip address 192.168.32.1 255.255.255.0
!
uplink wired priority 202
uplink cellular priority 201
uplink wired vlan  16
interface tunnel 2003
        description "Tunnel Interface"
        ip address 2.0.0.3 255.0.0.0
        tunnel source 192.168.30.1
        tunnel destination 192.168.68.217
        trusted
        ip ospf area 10.10.10.10
!
ip default-gateway 192.168.16.254
ip route 192.168.0.0 255.255.0.0 null 0
!
router ospf
router ospf router-id 192.168.30.1
router ospf area 10.10.10.10 stub
router ospf redistribute vlan 30-32
```

Remote Branch 2

```
controller-ip vlan 50
!
vlan 20
vlan 50
vlan 51
vlan 52
!
interface gigabitethernet  1/0
        description "GE1/0"
        trusted
        switchport access vlan 20
!
interface gigabitethernet  1/1
        description "GE1/1"
        trusted
        switchport access vlan 50
!
interface gigabitethernet  1/2
        description "GE1/2"
        trusted
        switchport access vlan 51
!
interface gigabitethernet  1/3
        description "GE1/3"
        trusted
        switchport access vlan 52
!
interface vlan 20
        ip address 192.168.20.1 255.255.255.0
!
interface vlan 50
```

```
        ip address 192.168.50.1 255.255.255.0
!
interface vlan 51
        ip address 192.168.51.1 255.255.255.0
!
interface vlan 52
        ip address 192.168.52.1 255.255.255.0
!
uplink wired priority 206
uplink cellular priority 205
uplink wired vlan  20
interface tunnel 2005
        description "Tunnel Interface"
        ip address 2.0.0.5 255.0.0.0
        tunnel source 192.168.50.1
        tunnel destination 192.168.68.217
        trusted
        ip ospf area 10.10.10.10
!
ip default-gateway 192.168.20.254
ip route 192.168.0.0 255.255.0.0 null 0
!
router ospf
router ospf router-id 192.168.50.1
router ospf area 10.10.10.10 stub
router ospf redistribute vlan 50-52
```

W-3200 Central Office Controller—Active

```
localip 0.0.0.0 ipsec db947e8d1b383813a4070ab0799fa6246b80fc5cfcc3268f
controller-ip vlan 225
vlan 68
vlan 100
vlan 225
!
interface gigabitethernet  1/0
        description "GE1/0"
        trusted
        switchport access vlan 225
!
interface gigabitethernet  1/1
        description "GE1/1"
        trusted
        switchport access vlan 100
!
interface gigabitethernet  1/2
        description "GE1/2"
        trusted
        switchport access vlan 68
!
interface vlan 68
        ip address 192.168.68.220 255.255.255.0
!
interface vlan 100
        ip address 192.168.100.1 255.255.255.0
!
interface vlan 225
        ip address 192.168.225.2 255.255.255.0
!
interface tunnel 2003
        description "Tunnel Interface"
        ip address 2.1.0.3 255.0.0.0
```

```
        tunnel source 192.168.225.2
        tunnel destination 192.168.30.1
        trusted
        ip ospf area 10.10.10.10
!
interface tunnel 2005
        description "Tunnel Interface"
        ip address 2.1.0.5 255.0.0.0
        tunnel source 192.168.225.2
        tunnel destination 192.168.50.1
        trusted
        ip ospf area 10.10.10.10
!
master-redundancy
  master-vrrp 2
  peer-ip-address 192.168.68.221 ipsec password123
!
vrrp 1
  priority 120
  authentication password123
  ip address 192.168.68.217
  vlan 68
  preempt
  tracking vlan 68 sub 40
  tracking vlan 100 sub 40
  tracking vlan 225 sub 40
  no shutdown
!
vrrp 2
  priority 120
  ip address 192.168.225.9
  vlan 225
  preempt
  tracking vlan 68 sub 40
  tracking vlan 100 sub 40
  tracking vlan 225 sub 40
  no shutdown
!
ip default-gateway 192.168.68.1
ip route 192.168.0.0 255.255.0.0 null 0

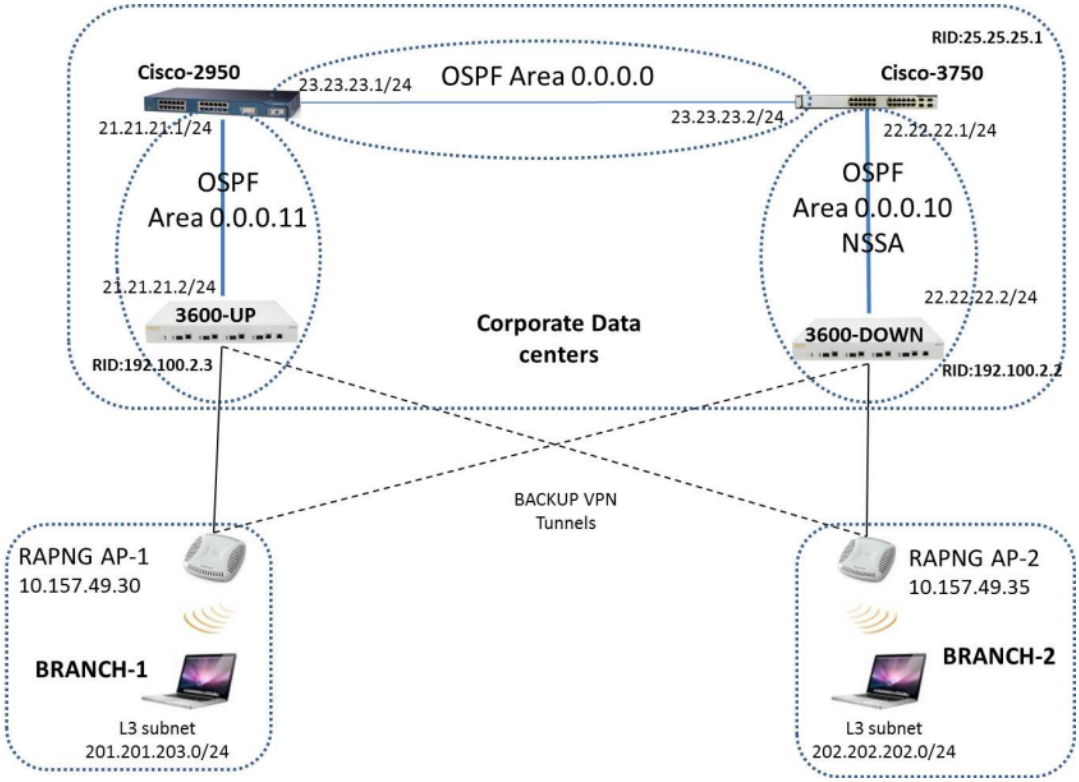
router ospf
router ospf router-id 192.168.225.1
router ospf area 10.10.10.10 stub
router ospf redistribute vlan 100,225
!
```

W-3200 Central Office Controller—Backup

```
localip 0.0.0.0 ipsec db947e8d1b383813a4070ab0799fa6246b80fc5cfcc3268f
controller-ip vlan 225
!
interface gigabitethernet  1/0
        description "GE1/0"
        trusted
        switchport access vlan 225
!
interface gigabitethernet  1/1
        description "GE1/1"
        trusted
        switchport access vlan 100
!
```

```
interface gigabitethernet  1/2
        description "GE1/2"
        trusted
        switchport access vlan 68
!
interface vlan 68
        ip address 192.168.68.221 255.255.255.224
!
interface vlan 100
        ip address 192.168.100.5 255.255.255.0
!
interface vlan 225
        ip address 192.168.225.1 255.255.255.0
!
interface tunnel 2003
        description "Tunnel Interface"
        ip address 2.1.0.3 255.0.0.0
        tunnel source 192.168.225.1
        tunnel destination 192.168.30.1
        trusted
        ip ospf area 10.10.10.10
!
interface tunnel 2005
        description "Tunnel Interface"
        ip address 2.1.0.5 255.0.0.0
        tunnel source 192.168.225.1
        tunnel destination 192.168.50.1
        trusted
        ip ospf area 10.10.10.10
!
master-redundancy
  master-vrrp 2
  peer-ip-address 192.168.68.220 ipsec password123
!
vrrp 1
  priority 99
  authentication password123
  ip address 192.168.68.217
  vlan 68
  tracking vlan 68 sub 40
  tracking vlan 100 sub 40
  tracking vlan 225 sub 40
  no shutdown
!
vrrp 2
  priority 99
  ip address 192.168.225.9
  vlan 225
  tracking vlan 68 sub 40
  tracking vlan 100 sub 40
  tracking vlan 225 sub 40
  no shutdown
!
ip default-gateway 192.168.68.1
ip route 192.168.0.0 255.255.0.0 null 0
!
router ospf
router ospf router-id 192.168.225.1
router ospf area 10.10.10.10 stub
router ospf redistribute vlan 100,225
!
```

The following figure displays how the controller is configured for Instant AP VPN for different OSPF cases.



Topology

- Area-10 is NSSA (Not-So-Stubby Area)
- Area-11 is Normal area.
- RAPNG AP-1 is configured to have a 3600-UP controller as its primary controller and a 3600-DOWN as secondary controller.
- RAPNG AP-2 is configured to have a 3600-DOWN as its primary controller and a 3600-UP as secondary controller.
- RAPNG AP-1 is configured to have a 201.201.203.0/24 L3-distributed network.
- RAPNG AP-2 is configured to have a 202.202.202.0/24 L3-distributed network.

Observation

- W-3600-UP Controller will send Type-5 LSA (External LSA) of VPN route 201.201.203.0/24 to it's upstream router, Cisco-3750.
- W-3600-DOWN Controller will send Type-7 LSA (NSSA) of VPN route 202.202.202.0/24 to it's upstream router, Cisco-2950.
- W-3600-UP Controller will send a Type-4 asbr-summary LSA.

Configuring W-3600-UP Controller

```
interface vlan 21
ip address 21.21.21.2 255.255.255.0
ip ospf area 0.0.0.11
!
router ospf
router ospf area 0.0.0.11
router ospf redistribute rapng-vpn
!
```

The following commands displays the configuration and run time protocol details on W-3600-UP Controller:

```
(host)#show ip route
Codes: C - connected, O - OSPF, R - RIP, S - static
M - mgmt, U - route usable, * - candidate default, V - RAPNG VPN
Gateway of last resort is Imported from DHCP to network 0.0.0.0 at cost 10
Gateway of last resort is Imported from CELL to network 0.0.0.0 at cost 10
Gateway of last resort is Imported from PPPOE to network 0.0.0.0 at cost 10
Gateway of last resort is 10.15.231.185 to network 0.0.0.0 at cost 1
S*    0.0.0.0/0  [1/0] via 10.15.231.185*
O    10.15.228.0/27 [333/0] via 21.21.21.1*
O    12.12.12.0/25 [0/0] via 21.21.21.1*
O    22.22.22.0/24 [3/0] via 21.21.21.1*
O    23.23.23.0/24 [2/0] via 21.21.21.1*
O    25.25.25.0/24 [333/0] via 21.21.21.1*
S    192.100.3.0/24 [1/0] via 192.100.2.1*
S    192.100.4.0/24 [1/0] via 192.100.2.1*
S    192.100.5.0/24 [1/0] via 192.100.2.1*
S    192.100.6.0/24 [1/0] via 192.100.2.1*
S    192.100.7.0/24 [1/0] via 192.100.2.1*
S    192.100.8.0/24 [1/0] via 192.100.2.1*
S    192.100.9.0/24 [1/0] via 192.100.2.1*
S    192.100.10.0/24 [1/0] via 192.100.2.1*
S    192.100.11.0/24 [1/0] via 192.100.2.1*
S    192.100.12.0/24 [1/0] via 192.100.2.1*
S    192.100.13.0/24 [1/0] via 192.100.2.1*
S    192.100.14.0/24 [1/0] via 192.100.2.1*
S    192.168.1.0/24 [1/0] via 192.100.2.1*
S    192.169.1.0/24 [1/0] via 192.100.2.1*
S    192.170.1.0/24 [1/0] via 192.100.2.1*
S    192.171.1.0/24 [1/0] via 192.100.2.1*
S    192.172.1.0/24 [1/0] via 192.100.2.1*
S    192.173.1.0/24 [1/0] via 192.100.2.1*
S    192.174.1.0/24 [1/0] via 192.100.2.1*
S    192.175.1.0/24 [1/0] via 192.100.2.1*
S    192.176.1.0/24 [1/0] via 192.100.2.1*
S    192.177.1.0/24 [1/0] via 192.100.2.1*
S    192.178.1.0/24 [1/0] via 192.100.2.1*
S    192.179.1.0/24 [1/0] via 192.100.2.1*
V    201.201.203.0/26 [10/0] ipsec map
O    202.202.202.0/29 [0/0] via 21.21.21.1*
C    192.100.2.0/24 is directly connected, VLAN2
C    10.15.231.184/29 is directly connected, VLAN1
C    172.16.0.0/24 is directly connected, VLAN3
C    21.21.21.0/24 is directly connected, VLAN21
C    5.5.0.2/32 is an ipsec map 10.15.149.30-5.5.0.2
```

```
(host) #show ip ospf database
OSPF Database Table
-------------------
Area ID    LSA Type        Link ID         Adv Router      Age    Seq#         Checksum
-------    --------        -------         ----------      ---    ----         --------
0.0.0.11   ROUTER          21.21.21.1      21.21.21.1      178    0x80000017   0xca50
0.0.0.11   ROUTER          192.100.2.3     192.100.2.3     1406   0x80000007   0x2253
0.0.0.11   NETWORK         21.21.21.1      21.21.21.1      178    0x80000003   0xdf6d
0.0.0.11   IPNET_SUMMARY   22.22.22.0      21.21.21.1      178    0x80000003   0x7e38
0.0.0.11   IPNET_SUMMARY   23.23.23.0      21.21.21.1      178    0x80000003   0x5064
0.0.0.11   ASBR_SUMMARY    25.25.25.1      21.21.21.1      178    0x80000003   0xefbc
0.0.0.11   ASBR_SUMMARY    192.100.2.3     192.100.2.3     1412   0x80000002   0xa85d
N/A        AS_EXTERNAL     10.15.228.0     25.25.25.1      1014   0x8000000e   0xea43
N/A        AS_EXTERNAL     12.12.12.0      25.25.25.1      268    0x80000003   0x433a
N/A        AS_EXTERNAL     25.25.25.0      25.25.25.1      1761   0x80000005   0x3d8d
```

```
N/A         AS_EXTERNAL     201.201.203.0   10.15.231.186   3600    0x80000001   0x6690
N/A         AS_EXTERNAL     201.201.203.0   192.100.2.3     1104    0x80000002   0xe4a2
N/A         AS_EXTERNAL     202.202.202.0   25.25.25.1      268     0x80000003   0x4385
```

```
(host) #show ip ospf neighbor
OSPF Neighbor Table
-------------------
Neighbor ID  Pri  State     Address      Interface
-----------  ---  -----     -------      ---------
21.21.21.1   1    FULL/DR   21.21.21.1   Vlan
```

Configuring W-3600-DOWN Controller

```
interface vlan 22
ip address 22.22.22.2 255.255.255.0
ip ospf area 0.0.0.10
!
router ospf
router ospf area 0.0.0.10 nssa
router ospf redistribute rapng-vpn
!
```

The following commands displays the configuration and run time protocol details on W-3600-DOWN Controller:

```
(host)#show ip route
Codes: C - connected, O - OSPF, R - RIP, S - static
M - mgmt, U - route usable, * - candidate default, V - RAPNG VPN
Gateway of last resort is Imported from DHCP to network 0.0.0.0 at cost 10
Gateway of last resort is Imported from CELL to network 0.0.0.0 at cost 10
Gateway of last resort is Imported from PPPOE to network 0.0.0.0 at cost 10
O    0.0.0.0/0  [1/0] via 22.22.22.1*
S    10.0.0.0/8  [1/0] via 10.15.231.177*
O    10.15.228.0/27 [333/0] via 22.22.22.1*
V    12.12.12.0/25 [10/0] ipsec map
O    21.21.21.0/24 [3/0] via 22.22.22.1*
O    23.23.23.0/24 [2/0] via 22.22.22.1*
O    25.25.25.0/24 [333/0] via 22.22.22.1*
V    202.202.202.0/29 [10/0] ipsec map
C    192.100.2.0/24 is directly connected, VLAN2
C    10.15.231.176/29 is directly connected, VLAN1
C    22.22.22.0/24 is directly connected, VLAN22
C    4.4.0.2/32 is an ipsec map 10.15.149.35-4.4.0.2
C    4.4.0.1/32 is an ipsec map 10.17.87.126-4.4.0.1
```

```
(host) #show ip ospf neighbor
OSPF Neighbor Table
-------------------
Neighbor ID  Pri  State      Address      Interface
-----------  ---  -----      -------      ---------
25.25.25.1   1    FULL/BDR   22.22.22.1   Vlan 22
```

```
(host) #show ip ospf database
OSPF Database Table
-------------------
Area ID    LSA Type        Link ID         Adv Router     Age    Seq#         Checksum
-------    --------        -------         ----------     ---    ----         --------
0.0.0.10   ROUTER          25.25.25.1      25.25.25.1     1736   0x80000021   0xb732
0.0.0.10   ROUTER          192.100.2.2     192.100.2.2    500    0x80000005   0x9ad9
0.0.0.10   NETWORK         22.22.22.2      192.100.2.2    500    0x80000004   0x8aeb
0.0.0.10   IPNET_SUMMARY   21.21.21.0      25.25.25.1     1990   0x80000003   0xe7bf
0.0.0.10   IPNET_SUMMARY   23.23.23.0      25.25.25.1     1990   0x80000003   0x950d
```

```
0.0.0.10    NSSA            0.0.0.0         25.25.25.1    725    0x80000002   0xaab9
0.0.0.10    NSSA            10.15.228.0     25.25.25.1    1228   0x80000010   0xca5f
0.0.0.10    NSSA            12.12.12.0      192.100.2.2   352    0x80000005   0xe8cb
0.0.0.10    NSSA            25.25.25.0      25.25.25.1    1485   0x80000006   0x1fa8
0.0.0.10    NSSA            202.202.202.0   192.100.2.2   352    0x80000005   0xe817
N/A         AS_EXTERNAL     12.12.12.0      192.100.2.2   352    0x80000005   0x28d8
N/A         AS_EXTERNAL     202.202.202.0   192.100.2.2   352    0x80000005   0x2824
```

Viewing the Status of Instant AP VPN

RAPNG AP-1

```
(host)# show  vpn status
profile name:default
-------------------------------------------------
current using tunnel                              :primary tunnel
ipsec is preempt status                           :disable
ipsec is fast failover status                     :disable
ipsec hold on period                              :600
ipsec tunnel monitor frequency (seconds/packet) :5
ipsec tunnel monitor timeout by lost packet cnt :2
ipsec     primary tunnel crypto type              :Cert
ipsec     primary tunnel peer address             :10.15.231.186
ipsec     primary tunnel peer tunnel ip           :192.100.2.3
ipsec     primary tunnel ap tunnel ip             :5.5.0.2
ipsec     primary tunnel current sm status        :Up
ipsec     primary tunnel tunnel status            :Up
ipsec     primary tunnel tunnel retry times       :2
ipsec     primary tunnel tunnel uptime            :1 hour 24 minutes 50 seconds
ipsec      backup tunnel crypto type              :Cert
ipsec      backup tunnel peer address             :10.15.231.178
ipsec      backup tunnel peer tunnel ip           :0.0.0.0
ipsec      backup tunnel ap tunnel ip             :0.0.0.0
ipsec      backup tunnel current sm status        :Init
ipsec      backup tunnel tunnel status            :Down
ipsec      backup tunnel tunnel retry times       :0
ipsec      backup tunnel tunnel uptime            :0
(host)# show datapath route
Route Table Entries
-------------------
Flags: L - Local, P - Permanent,  T - Tunnel, I - IPsec, M - Mobile, A - ARP, D - Drop
IP              Mask            Gateway         Cost  VLAN  Flags
--------------  --------------  --------------  ----  ----  -----
0.0.0.0         0.0.0.0         10.15.149.25       0     0
0.0.0.0         128.0.0.0       192.100.2.3        0     0  T
128.0.0.0       128.0.0.0       192.100.2.3        0     0  T
192.168.10.0    255.255.254.0   192.168.10.1       0  3333  D
201.201.203.0   255.255.255.192 0.0.0.0            0   103  LP
10.15.149.24    255.255.255.248 10.15.149.30       0     1  L
10.15.231.186   255.255.255.255 10.15.149.25       0     0
Route Cache Entries
-------------------
Flags: L - local, P - Permanent,  T - Tunnel, I - IPsec, M - Mobile, A - ARP, D - Drop
IP              MAC               VLAN      Flags
--------------  ----------------  ----------  -----
202.202.202.6   00:00:00:00:00:00           0  T
192.100.2.3     00:00:00:00:00:00           0  PT
192.168.10.51   10:40:F3:98:80:94           1  PA
192.168.10.1    00:24:6C:C9:27:A3        3333  LP
201.201.203.8   00:26:C6:52:6B:14         103
201.201.203.1   00:24:6C:C9:27:A3         103  LP
10.1.1.50       00:00:00:00:00:00           0  T
```

```
5.5.0.2          00:24:6C:C9:27:A3            1  LP
10.15.149.30     00:24:6C:C9:27:A3            1  LP
10.15.149.25     00:0B:86:40:93:00            1  A
(host)# show clients
Client List
-----------
Name  IP Address     MAC Address        OS  Network  Access Point       Channel  Type  Role
Signal    Speed (mbps)
----  ----------     -----------        --  -------  ------------       -------  ----  ----
------    ------------
201.201.203.8  00:26:c6:52:6b:14      149.30   00:24:6c:c9:27:a3  48-      AN    149.30  43
(good)  6(poor)
Info timestamp      :80259
```

RAPNG AP-3

```
(host)# show vpn status
profile name:default
-------------------------------------------------
current using tunnel                            :primary tunnel
ipsec is preempt status                         :disable
ipsec is fast failover status                   :disable
ipsec hold on period                            :600
ipsec tunnel monitor frequency (seconds/packet) :5
ipsec tunnel monitor timeout by lost packet cnt :2
ipsec     primary tunnel crypto type            :Cert
ipsec     primary tunnel peer address           :10.15.231.178
ipsec     primary tunnel peer tunnel ip         :192.100.2.2
ipsec     primary tunnel ap tunnel ip           :4.4.0.2
ipsec     primary tunnel current sm status      :Up
ipsec     primary tunnel tunnel status          :Up
ipsec     primary tunnel tunnel retry times     :13
ipsec     primary tunnel tunnel uptime          :1 hour 55 minutes 6 seconds
ipsec      backup tunnel crypto type            :Cert
ipsec      backup tunnel peer address           :10.15.231.186
ipsec      backup tunnel peer tunnel ip         :0.0.0.0
ipsec      backup tunnel ap tunnel ip           :0.0.0.0
ipsec      backup tunnel current sm status      :Init
ipsec      backup tunnel tunnel status          :Down
ipsec      backup tunnel tunnel retry times     :0
ipsec      backup tunnel tunnel uptime          :0
```

```
(host)# show datapath route
Route Table Entries
-------------------
Flags: L - Local, P - Permanent,  T - Tunnel, I - IPsec, M - Mobile, A - ARP, D - Drop
IP              Mask            Gateway         Cost  VLAN  Flags
--------------  --------------  --------------  ----  ----  -----
0.0.0.0         0.0.0.0         10.15.149.33       0     0
0.0.0.0         128.0.0.0       192.100.2.2        0     0  T
128.0.0.0       128.0.0.0       192.100.2.2        0     0  T
192.168.10.0    255.255.254.0   192.168.10.1       0  3333  D
10.15.149.32    255.255.255.248 10.15.149.35       0     1  L
202.202.202.0   255.255.255.248 0.0.0.0            0   203  LP
10.15.231.178   255.255.255.255 10.15.149.33       0     0
Route Cache Entries
-------------------
Flags: L - local, P - Permanent,  T - Tunnel, I - IPsec, M - Mobile, A - ARP, D - Drop
IP              MAC             VLAN      Flags
--------------  --------------  ----------  -----
202.202.202.1   00:24:6C:C0:41:F2          203  LP
202.202.202.6   08:ED:B9:E1:51:7B          203
```

```
192.100.2.2       00:00:00:00:00:00            0  PT
192.168.10.1      00:24:6C:C0:41:F2         3333  LP
201.201.203.8     00:00:00:00:00:00            0  T
10.1.1.50         00:00:00:00:00:00            0  T
192.168.11.7      00:26:C6:52:6B:14            1  PA
4.4.0.2           00:24:6C:C0:41:F2            1  LP
10.13.6.110       00:00:00:00:00:00            0  T
10.15.149.38      00:24:6C:C9:27:CC            1  A
10.15.149.35      00:24:6C:C0:41:F2            1  LP
10.15.149.33      00:0B:86:40:93:00            1  A
```

```
(host)# show clients
Client List
-----------
Name  IP Address     MAC Address        OS  Network  Access Point       Channel  Type  Role
Signal    Speed (mbps)
----  -----------    -----------        --  -------  ------------       -------  ----  ----
------    ------------
202.202.202.6  08:ed:b9:e1:51:7b      149.35   00:24:6c:c0:41:f2  48-      AN    149.35  53
(good)  48(poor)
Info timestamp      :80748
```

This chapter describes how to configure a Dell tunneled node, also known as a wired tunneled node. Dell tunneled nodes provide access and security using an overlay architecture.

This chapter describes the following topics:

- Understanding Tunneled Node Configuration on page 222
- Configuring a Wired Tunneled Node Client on page 223

Understanding Tunneled Node Configuration

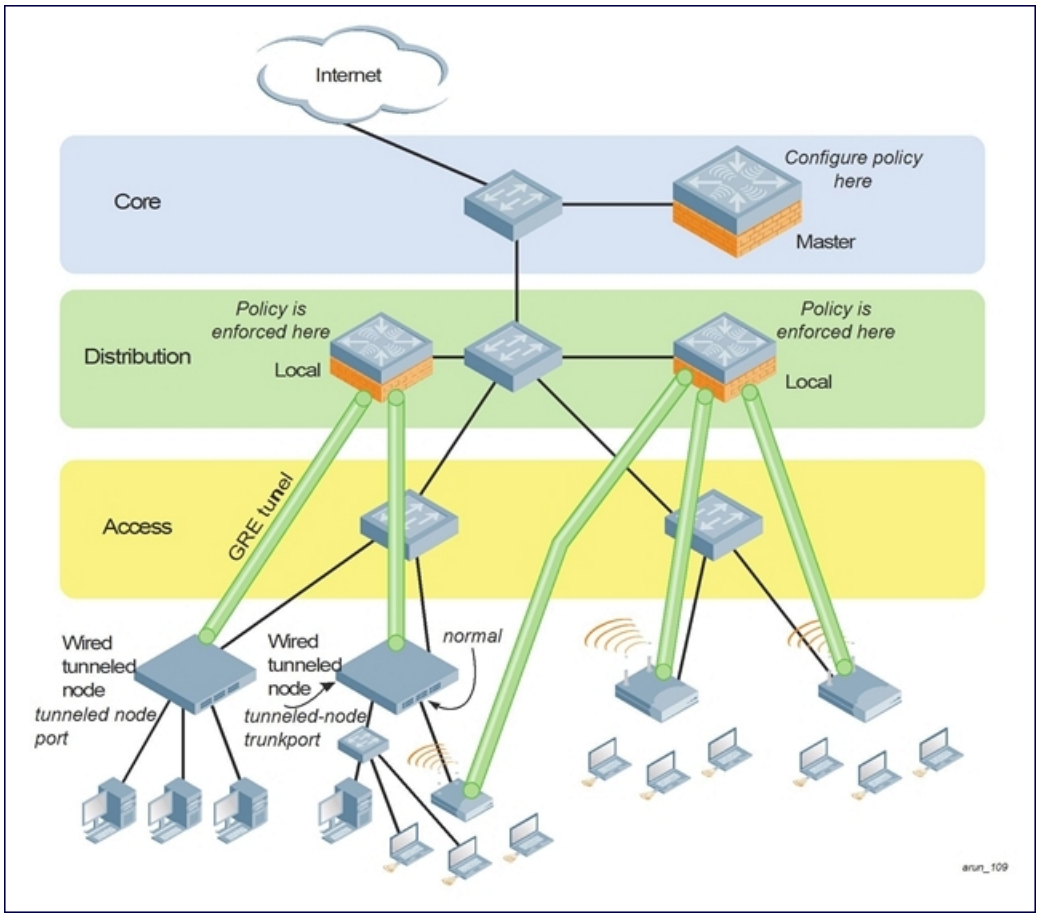
The Dell tunneled node connects to one or more client devices at the edge of the network and then establishes a secure GRE tunnel to the controlling concentrator server. This approach allows the controller to support all the centralized security features, such as 802.1x authentication, captive-portal authentication, and stateful firewall. The Dell tunneled node is required to handle only the physical connection to clients and support for its end of the GRE tunnel.

To support the wired concentrator, the controller must have a license to terminate access points (APs). No other configuration is required. To configure the Dell tunneled node, you must specify the IP address of the controller and identify the ports that are to be used as active tunneled node ports. Tunnels are established between the controller and each active tunneled node port on the tunneled node. All tunneled node units must be running the same version of software. The tunneled node port can also be configured as a trunk port. This allows customers to have multiple clients on different VLANs that come through the trunk port instead of having clients on a single vlan.

Figure 27 shows how the tunneled node fits into network operations. Traffic moves through GRE tunnels between the active tunneled node ports and the controller or controllers. Policies are configured on a master server and enforced on the local controllers. The master and the local controller can run on the same or different systems. The tunneled node can connect to the master, but it is not required.

On the controlling controller, you can assign the same policy to tunneled node user traffic as you would to any untrusted wired traffic. The profile specified by the **`aaa authentication wired`** command determines the initial role, which contains the policy. The VLAN setting on the concentrator port must match the VLAN that will be used for users at the local controller.

Figure 27 *Tunneled Node Configuration Operation*



Configuring a Wired Tunneled Node Client

NOTE

ArubaOS does not allow a tunneled-node client and tunneled-node server to co-exist on the same controller at the same time. The controller must be configured as either a tunneled-node client or a tunneled-node server. By default, the controller behaves as a tunneled-node server. However, once tunneled-node-server xxx.xxx.xxx.xxx is configured on the controller, the controller becomes a tunneled-node client. To remove the tunneled-node client function, use the command **tunneled-node-server 0.0.0.0** to disable the tunneled-node client on the controller side.

This section describes how to configure a tunneled node client. You can use the WebUI or the CLI to complete the configuration steps.

1. Access the Wired tunneled node CLI according to the instructions provided in the installation guide that shipped with your tunneled node. Console access (9600 8N1) and SSH access are supported.
2. Specify the IP address of the controller and specify tunnel loop prevention.
   - CLI:

```
(host)(config) #tunneled-node-address ipaddress
(host)(config) #tunnel-loop-prevention
```

   - WebUI

   a. Navigate to **Configuration>Advanced Services>Wired Access** page.
   b. Locate the **Wired Access Concentration Configuration** section.
   c. T o enable tunneled nodes, click the **Enable Wired Access Concentrator** checkbox.

d. Enter the IP address of the controller in the **Wired Access Concentrator Server IP** field.
e. To enable tunnel loop prevention, click the **Enable Wired Access Concentrator Loop Prevention** checkbox.
f. Click **Apply**.

3. Access each interface that you want to use, and assign it as a tunneled node port.

```
(host (config) # interface fastethernet n/m
(host (config-if) # tunneled-node port
```

4. Verify the configuration.

```
(host) (config-if) # exit
(host) # show tunneled-port config
```

Configuring an Access Port as a Tunneled Node Port

You can configure any port on any controller as a tunneled node port using the tunneled-node-port command. Set the tunneled-nod -address as the controller to act as the tunneled node termination point. The **tunneled-node-port** command tells the physical interface to tunnel that traffic to the controller.

1. Enable portfast on the Wired tunneled node.

```
(host)(config) #interface fastethernet <slot>/<port>
(host) (config-if) # spanning-tree portfast
```

2. Assign a VLAN to the tunneled node port.

```
(host) (config-if) # switchport mode access
(host) (config-if) # switchport access vlan <vlanid>
```

Configuring a Trunk Port as a Tunneled Node Port

To enable portfast on the Wired tunneled node execute the following commands:

- `(host) (config-if) # switchport mode trunk`
- `(host) (config-if) # switchport trunk allowed vlan <WORD>`

Show commands

To verify the status of the Wired tunneled node execute the following command:

- `(host) # show tunneled-node state`
- `(host) #show tunneled-node config`

To check the current usage on the controller execute the following command. Each tunneled node client uses one AP license. Attaching an additional wired client on the tunneled node client does not increment the AP license usage on the controller.

- `(host) #show license-usage ap`

The ArubaOS software allows you to use an external authentication server or the controller internal user database to authenticate clients who need to access the wireless network.

This chapter describes the following topics:

- Understanding Authentication Server Best Practices and Exceptions on page 225
- Understanding Servers and Server Groups on page 225
- Configuring Authentication Servers on page 226
- Managing the Internal Database on page 235
- Configuring Server Groups on page 238
- Assigning Server Groups on page 244
- Configuring Authentication Timers on page 248
- Authentication Server Load Balancing on page 249

Understanding Authentication Server Best Practices and Exceptions

- For an external authentication server to process requests from the Dell controller, you must configure the server to recognize the controller. Refer to the vendor documentation for information on configuring the authentication server.
- Instructions on how to configure Microsoft's IAS and Active Directory can be viewed at:

  Microsoft's IAS:

  technet2.microsoft.com/windowsserver/en/technologies/ias.mspx

  Active Directory:

  microsoft.com/en-us/server-cloud/windows-server/active-directory.aspx

Understanding Servers and Server Groups

ArubaOS supports the following external authentication servers:

- RADIUS (Remote Authentication Dial-In User Service)
- LDAP (Lightweight Directory Access Protocol)
- TACACS+ (Terminal Access Controller Access Control System)
- Windows (For stateful NTLM authentication)



Starting from ArubaOS 6.4, a maximum of 128 LDAP, RADIUS, and TACACS servers, each can be configured on the controller.

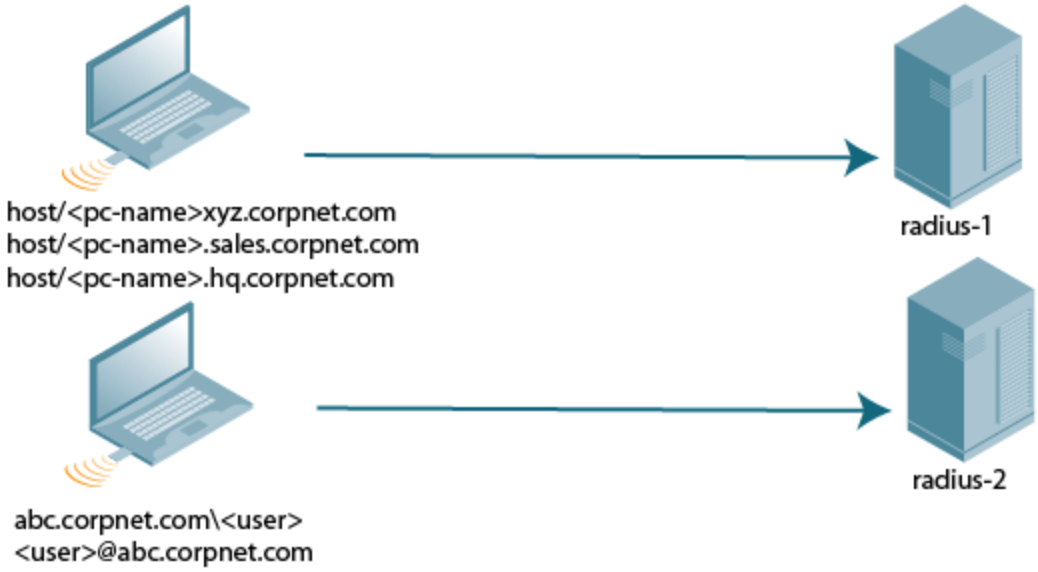
Additionally, you can use the controller's internal database to authenticate users. You create entries in the database for users, their passwords, and their default role.

You can create *groups* of servers for specific types of authentication. For example, you can specify one or more RADIUS servers to be used for 802.1x authentication. The list of servers in a server group is an ordered list. This means that the first server in the list is always used unless it is unavailable, in which case the next server in

the list is used. You can configure servers of different types in one group. For example, you can include the internal database as a backup to a RADIUS server.

Figure 28 graphically represents a server group named "Radii" that consists of two RADIUS servers, Radius-1 and Radius-2. The server group is assigned to the server group for 802.1x authentication.

Figure 28 *Server Group*



Server names are unique. You can configure the same server in multiple server groups. You must configure the server before you can add it to a server group.

NOTE

If you use the controller's internal database for user authentication, use the predefined "Internal" server group.

You can also include conditions for server-derived user roles or VLANs in the server group configuration. The server derivation rules apply to all servers in the group.

Configuring Authentication Servers

This section describes how to configure RADIUS, LDAP, TACACS+ and Windows external authentication servers and the internal database on the controller.

This section includes the following information:

- Configuring a RADIUS Server on page 226
- Configuring an RFC-3576 RADIUS Server on page 232
- Configuring an LDAP Server on page 232
- Configuring a TACACS+ Server on page 234
- Configuring a Windows Server on page 235

Configuring a RADIUS Server

Follow the procedures below to configure a RADIUS server using the WebUI or CLI.

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.

2. Select **Radius Server** to display the Radius Server List.
3. To configure a RADIUS server, enter the name for the server and click **Add.**
4. Select the name to configure server parameters. Enter parameters as described in Table 36. Select the **Mode** checkbox to activate the authentication server.
5. Click **Apply**.



The configuration does not take effect until you perform this step.

Using the CLI

```
(host)(config) #aaa authentication-server radius <name>
  host <ipaddr>
  key <key>
  enable
```

Table 36: *RADIUS Server Configuration Parameters*

Parameter	Description
Host	IP address or fully qualified domain name (FQDN) of the authentication server. The maximum supported FQDN length is 63 characters. Default: N/A
Key	Shared secret between the controller and the authentication server. The maximum length is 128 characters. Default: N/A
Authentication Port	Authentication port on the server. Default: 1812
Accounting Port	Accounting port on the server Default: 1813
Retransmits	Maximum number of retries sent to the server by the controller before the server is marked as down. Default: 3
Timeout	Maximum time, in seconds, that the controller waits before timing out the request and resending it. Default: 5 seconds
NAS ID	Network Access Server (NAS) identifier to use in RADIUS packets. Default: N/A
NAS IP	NAS IP address to send in RADIUS packets. You can configure a "global" NAS IP address that the controller uses for communications with all RADIUS servers. If you do not configure a server-specific NAS IP, the global NAS IP is used. To set the global NAS IP in the WebUI, navigate to the **Configuration > Security > Authentication > Advanced** page. To set the global NAS IP in the CLI, enter the **ip radius nas-ip** <*ipaddr*> command. Default: N/A
Source Interface	Enter a VLAN number ID. Allows you to use source IP addresses to differentiate RADIUS requests. Associates a VLAN interface with the RADIUS server to allow the server-specific source interface to override the global configuration. • If you associate a Source Interface (by entering a VLAN number) with a configured server, then the source IP address of the packet is that interface's IP address.

Parameter	Description
	• If you do not associate the Source Interface with a configured server (leave the field blank), the IP address of the global Source Interface is used.
Use MD5	Use MD5 hash of cleartext password. Default: Disabled
Mode	Enables or disables the server. Default: Enabled

RADIUS Server VSAs

Vendor-Specific Attributes (VSAs) are a method for communicating vendor-specific information between Network Access Servers and RADIUS servers, allowing vendors to support their own extended attributes. You can use DellVSAs to derive the user role and VLAN for RADIUS-authenticated clients; however the VSAs must be present on your RADIUS server. This requires that you update the RADIUS dictionary file with the vendor name (Aruba) and/or the vendor-specific code (14823), the vendor-assigned attribute number, and the attribute format (such as string or integer) for each VSA. For more information on VSA-derived user roles, see Configuring a VSA-Derived Role on page 376

The following table describes Dell-specific RADIUS VSAs. For the current and complete list of all RADIUS VSAs available in the version of ArubaOS currently running on your controller, access the command-line interface and issue the command **show aaa radius attributes**.

Table 37: *RADIUS VSAs*

VSA	Type	Value	Description
Aruba-User-Role	String	1	This VSA returns the role, to be assigned to the user post authentication. The user will be granted access based on the role attributes defined in the role.
Aruba-User-Vlan	Integer	2	This VSA returns the VLAN to be used by the client. Range: 1–4094.
Aruba-Priv-Admin-User	Integer	2	If this VSA is set in the RADIUS accept message, the user can bypass the enable prompt.
Aruba-Admin-Role	String	4	This VSA returns the management role to be assigned to the user post management authentication. This role can be seen using the command **show mgmt-role** in the command-line interface.
Aruba-Essid-Name	String	5	String that identifies the name of the ESSID.
Aruba-Location-Id	String	6	String that identifies the name of the AP location.
Aruba-Port-Id	String	7	String that identifies the Port ID.
Aruba-Template-User	String	8	String that identifies the name of a Dell user template.

VSA	Type	Value	Description
Aruba-Named-User-Vlan	String	9	This VSA returns a VLAN name for a user. This VLAN name on a controller could be mapped to user-defined name or or multiple VLAN IDs.
Aruba-AP-Group	String	10	String that identifies the name of a Dell AP Group.
Aruba-Framed-IPv6-Address	String	11	This attribute is used for RADIUS accounting for IPv6 users.
Aruba-Device-Type	String	12	String that identifies a Dell device on the network.
Aruba-No-DHCP-Fingerprint	Integer	14	This VSA prevents the controller from deriving a role and VLAN based on DHCP finger printing.
Aruba-Mdps-Device-Udid	String	15	UDID is unique device identifier which is used as input attribute by the Onboard application while performing the device authorization to the internal RADIUS server within the ClearPass Policy Manager (CPPM). The UDID checks against role mappings or enforcement policies to determine if the device is authorized to be onboarded.
Aruba-Mdps-Device-Imei	String	16	The Onboard application uses IMEI as an input attribute while performing the device authorization to the internal RADIUS server within the CPPM. IMEI checks against role mappings or enforcement policies to determine if the device is authorized to be onboarded.
Aruba-Mdps-Device-Iccid	String	17	The Onboard application uses ICCID as an input attribute while performing the device authorization to the internal RADIUS server within the CPPM. ICCID checks against role mappings or enforcement policies to determine if the device is authorized to be onboarded.
Aruba-Mdps-Max-Devices	String	18	Used by Onboard as a way to define and enforce the maximum number of devices that can be provisioned by a given user.
Aruba-Mdps-Device-Name	String	19	The Onboard application uses device name as an input attribute while performing the device authorization to the internal RADIUS server within the CPPM. Device name checks against role mappings or enforcement policies to determine if the device is authorized to be onboarded.
Aruba-Mdps-Device-Product	String	20	The device product is used as input attribute by the Onboard application while performing the device authorization to the internal RADIUS server within the CPPM. Device Product checks against role mappings or enforcement policies to determine if the device is authorized to be onboarded.

VSA	Type	Value	Description
Aruba-Mdps-Device-Version	String	21	The device version is used as input attribute by the Onboard application while performing the device authorization to the internal RADIUS server within the CPPM. Device Version checks against role mappings or enforcement policies to determine if the device is authorized to be onboarded.
Aruba-Mdps-Device-Serial	String	22	The device serial number is used as input attribute by the Onboard application while performing the device authorization to the internal RADIUS server within the CPPM. Device Serial checks against role mappings or enforcement policies to determine if the device is authorized to be onboarded.
Aruba-AirGroup-User-Name	String	24	A device owner or username associated with the device.
Aruba-AirGroup-Shared-User	String	25	This VSA contains a comma separated list of user names with whom the device is shared.
Aruba-AirGroup-Shared-Role	String	26	This VSA contains a comma separated list of user roles with whom the device is shared.
Aruba-AirGroup-Device-Type	Integer	27	A value of 1 for this VSA indicates that the device authenticating on the network is a personal device and a value of 2 indicates that it is a shared device.
Aruba-Auth-Survivability	String	28	The Instant AP Auth survivability feature uses the VSA to indicate that the CPPM server sends the **Aruba-AS-User-Name** and **Aruba-AS-Credential-Hash** values. This attribute is just used as a flag with no specific value required.
Aruba-AS-User-Name	String	29	The Auth survivability feature uses the VSA for Instant APs. The CPPM sends the actual user name to the Instant AP which can be used by the Instant AP to authenticate the user if the CPPM server is not reachable.
Aruba-AS-Credential-Hash	String	30	The Auth survivability feature uses the VSA for Instant APs. The CPPM sends the NT hash of the password to the Instant AP which can be used by the Instant AP to authenticate the user if the CPPM server is not reachable.

VSA	Type	Value	Description
Aruba-WorkSpace-App-Name	String	31	This VSA identifies an application supported by Dell WorkSpace.
Aruba-Mdps-Provisioning-Settings	String	32	Used as part of the ClearPass Onboard technology, this attribute allows the CPPM to signal back to the onboard process the context of the device provisioning settings that should be applied to the device based on applied role mappings.
Aruba-Mdps-Device-Profile	String	33	Used as part of the ClearPass Onboard technology, this attribute allows CPPM to signal back to the onboard process the device profile that should be applied to the device based on applied role mappings.

RADIUS Server Authentication Codes

A configured RADIUS server returns the following standard response codes.

Table 38: *RADIUS Authentication Response Codes*

Code	Description
0	Authentication OK.
1	Authentication failed : user/password combination not correct.
2	Authentication request timed out : No response from server.
3	Internal authentication error.
4	Bad Response from RADIUS server : verify shared secret is correct.
5	No RADIUS authentication server is configured.
6	Challenge from server. (This does not necessarily indicate an error condition.)

RADIUS Server Fully Qualified Domain Names

If you define a RADIUS server using the FQDN of the server rather than its IP address, the controller periodically generates a DNS request and caches the IP address returned in the DNS response. To view the IP address that currently correlates to each RADIUS server FQDN, access the command-line interface in config mode and issue the following command:

```
show aaa fqdn-server-names
```

DNS Query Intervals

If you define a RADIUS server using the FQDN of the server rather than its IP address, the controller periodically generates a DNS request and caches the IP address returned in the DNS response. DNS requests are sent every 15 minutes by default.

You can use either the WebUI or the CLI to configure how often the controller will generate a DNS request to cache the IP address for a RADIUS server identified via its fully qualified domain name (FQDN).

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Advanced** page.
2. In the **DNS Query Interval** (min) field, enter a new DNS query interval, from 1-1440 minutes, inclusive.
3. Click **Apply**.

Using the CLI

```
(host)(config) #aaa dns-query-period <minutes>
```

Configuring an RFC-3576 RADIUS Server

You can configure a RADIUS server to send user disconnect, change-of-authorization (CoA), and session timeout messages as described in RFC 3576, "Dynamic Authorization Extensions to Remote Dial In User Service (RADIUS)."

The disconnect, session timeout, and change-of-authorization messages sent from the server to the controller contains information to identify the user for which the message is sent. The controller supports the following attributes for identifying the users who authenticate with an RFC 3576 server:

- user-name: name of the user to be authenticated
- framed-ip-address: user's IP address
- calling-station-id: phone number of a station that originated a call
- accounting-session-id: unique accounting ID for the user session.

If the authentication server sends both supported and unsupported attributes to the controller, the unknown or unsupported attributes are ignored. If no matching user is found, the controller sends a *503: Session Not Found* error message back to the RFC 3576 server.

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **RFC 3576 Server** to display the Radius Server List.
3. To define a new RFC 3576 RADIUS server, enter the IP address for the server and click **Add**.
4. Select the server name to configure server parameters.
5. Enter the server authentication key into the **Key** and **Retype** fields.
6. Click **Apply**.



The configuration does not take effect until you perform this step.

Using the CLI

```
(host)(config) #aaa rfc-3576-server <ipaddr>
  clone <server>
  key <psk>
  no ...
```

Configuring an LDAP Server

Table 39 describes the parameters you configure for an LDAP server.

Table 39: *LDAP Server Configuration Parameters*

Parameter	Description
Host	IP address of the LDAP server. Default: N/A
Admin-DN	Distinguished name for the admin user who has read/search privileges across all the entries in the LDAP database (the user does need write privileges, but will be able to search the database, and read attributes of other users in the database).
Admin Password	Password for the admin user. Default: N/A
Allow Clear-Text	Allows clear-text (unencrypted) communication with the LDAP server. Default: disabled
Authentication Port	Port number used for authentication. Default: 389
Base-DN	Distinguished Name of the node that contains the entire user database. Default: N/A
Filter	A string searches for users in the LDAP database. The default filter string is: **(objectclass=*)**. Default: N/A
Key Attribute	A string searches for a LDAP server. For Active Directory, the value is sAMAccountName. Default: sAMAccountName
Timeout	Timeout period of a LDAP request, in seconds. Default: 20 seconds
Mode	Enables or disables the server. Default: enabled
Preferred Connection Type	Preferred type of connection between the controller and the LDAP server. The default order of connection type is: 1. ldap-s 2. start-tls 3. clear-text The controller first tries to contact the LDAP server using the preferred connection type, and only attempts to use a lower-priority connection type if the first attempt is not successful. **NOTE:** If you select **clear-text** as the preferred connection type, you must also enable the **allow-cleartext** option.

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **LDAP Server** to display the LDAP Server List.
3. To configure an LDAP server, enter the name for the server and click **Add**.
4. Select the name to configure server parameters. Enter parameters as described in Table 39. Select the **Mode** checkbox to activate the authentication server.
5. Click **Apply**.



The configuration does not take effect until you perform this step.

Using the CLI

```
(host)(config) #aaa authentication-server ldap <name>
  host <ipaddr>
```

(enter parameters as described in Table 39)

```
  enable
```

Configuring a TACACS+ Server

Table 40 defines the TACACS+ server parameters.

Table 40: *TACACS+ Server Configuration Parameters*

Parameter	Description
Host	IP address of the server. Default: N/A
Key	Shared secret to authenticate communication between the TACACS+ client and server. Default: N/A
TCP Port	TCP port used by server. Default: 49
Retransmits	Maximum number of times a request is retried. Default: 3
Timeout	Timeout period for TACACS+ requests, in seconds. Default: 20 seconds
Mode	Enables or disables the server. Default: enabled
Session Authorization	Enables or disables session authorization. Session authorization turns on the optional authorization session for admin users. Default: disabled

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **TACACS Server** to display the TACACS Server List.
3. To configure a TACACS+ server, enter the name for the server and click **Add**.
4. Select the name to configure server parameters. Enter parameters as described in Table 40. Select the **Mode** checkbox to activate the authentication server.
5. Click **Apply**.



The configuration does not take effect until you perform this step.

Using the CLI

The following command configures, enables a TACACS+ server and enables session authorization:

```
(host)(config) #aaa authentication-server tacacs <name>
  clone default
  host <ipaddr>
  key <key>
  enable
```

```
session-authorization
```

Configuring a Windows Server

Table 41 defines parameters for a Windows server used for stateful NTLM authentication.

Table 41: *Windows Server Configuration Parameters*

Parameter	Description
Host	IP address of the server. Default: N/A
Mode	Enables or disables the server. Default: enabled
Windows Domain	Name of the Windows Domain assigned to the server.

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **Windows Server** to display the Windows Server List.
3. To configure a Windows server, enter the name for the server and click **Add**.
4. Select the name of the server to configure its parameters. Enter the parameters as described in Table 41.
5. Select the **Mode** checkbox to activate the authentication server.
6. Click **Apply**.



The configuration does not take effect until you perform this step.

Using the CLI

```
aaa authentication-server windows <windows-server-name>
   host <ipaddr>
   enable
```

Managing the Internal Database

You can create entries in the controller's internal database, to use to authenticate clients. The internal database contains a list of clients, along with the password and default role for each client. When you configure the internal database as an authentication server, client information is checked in incoming authentication requests against the internal database.

Configuring the Internal Database

The master controller uses the internal database for authentication by default. You can choose to use the internal database in a local controller by entering the CLI command **`aaa authentication-server internal use-local-switch`**. If you use the internal database in a local controller, you need to add clients on the local controller.

Table 42 defines the required and optional parameters used in the internal database.

Table 42: *Internal Database Configuration Parameters*

Parameters	Description
User Name	(Required) Enter a user name or select **Generate** to automatically generate a user name. An entered user name can be up to 64 characters in length.
Password	(Required) Enter a password or select **Generate** to automatically generate a password string. An entered password must be a minimum of 6 characters and can be up to 128 characters in length.
Role	Role for the client. For this role to be assigned to a client, you need to configure a server derivation rule, as described in Configuring Server-Derivation Rules on page 242. (A user role assigned through a server-derivation rule takes precedence over the default role configured for an authentication method.)
E-mail	(Optional) E-mail address of the client.
Enabled	Select this checkbox to enable the user as soon as the user entry is created.
Expiration	Select one of the following options: • Entry does not expire: No expiration on user entry. • Set Expiry time (mins): Enter the number of minutes the user is authenticated before their user entry expires. • Set Expiry Date (mm/dd/yyyy) Expiry Time (hh:mm): To select a specific expiration date and time, enter the expiration date in mm/dd/yyyy format, and the expiration time in hh:mm format.
Static Inner IP Address (for RAPs only)	Assign a static inner IP address to a Remote AP. If this database entry is not for a remote AP, leave this field empty.

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **Internal DB**.
3. Click **Add User** in the **Users** section. The user configuration page displays.
4. Enter the information for the client, as described in the table above.
5. Click **Enabled** to activate this entry on creation.
6. Click **Apply**. The configuration does not take effect until you perform this step.
7. At the Servers page, click **Apply**.



The Internal DB Maintenance window also includes a **Guest User Page** feature that allows you to create user entries for guests only. For details on creating guest users, see Guest Provisioning User Tasks on page 814.

Using the CLI

Enter the following command in enable mode:

```
(host)(config) #local-userdb add {generate-username|username <name>}{
generate-password|password <password>}
```

Managing Internal Database Files

ArubaOS allows you to import and export tables of user information to and from the internal database. These files should not be edited once they are exported. ArubaOS only supports the importing of database files that

were created during the export process. Note that importing a file into the internal database overwrite and removes all existing entries.

Exporting Files in the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **Internal DB**.
3. Click **Export** in the **Internal DB Maintenance** section. A popup window opens.
4. Enter the name of the file you want to export
5. Click **OK**.

Importing Files in the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **Internal DB**.
3. Click **Import** in the **Internal DB Maintenance** section. A popup window opens.
4. Enter the name of the file you want to import.
5. Click **OK**.

Exporting and Importing Files in the CLI

Enter the following command in enable mode:

```
(host)(config) #local-userdb export <filename>
(host)(config) #local-userdb import <filename>
```

Working with Internal Database Utilities

The local internal database also includes utilities to clear all users from the database and to restart the internal database to repair internal errors. Under normal circumstances, neither of these utilities are necessary.

Deleting All Users

Issue this command to remove users from the internal database after you have moved your user database from the controller's internal server to an external server.

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **Internal DB**.
3. Click **Delete All Users** in the **Internal DB Maintenance** section. A popup window open and asks you to confirm that you want to remove all users.
4. Click **OK**.

Repairing the Internal Database

Use this utility under the supervision of Dell technical support to recreate the internal database. This may clear internal database errors, but also removes all information from the database. Make sure you export your current user information before you start the repair procedure.

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **Internal DB**.
3. Click **Repair Database** in the **Internal DB Maintenance** section. A popup window open and asks you to confirm that you want to recreate the database.
4. Click **OK**.

Configuring Server Groups

You can create *groups* of servers for specific types of authentication – for example, you can specify one or more RADIUS servers to be used for 802.1x authentication. You can configure servers of different types in one group. For example, you can include the internal database as a backup to a RADIUS server.

Configuring Server Groups

Server names are unique. You can configure the same server in more than one server group. You must configure the server before you can include it in a server group.

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **Server Group** to display the Server Group list.
3. Enter the name of the new server group and click **Add**.
4. Select the name to configure the server group.
5. Under Servers, click **New** to add a server to the group.
   a. Select a server from the drop-down list and click **Add Server**.
   b. Repeat the above step to add other servers to the group.
6. Click **Apply**.

Using the CLI

```
(host)(config) #aaa server-group <name>
   auth-server <name>
```

Configuring Server List Order and Fail-Through

The list of servers in a server group is an ordered list. The first server in the list is always used by default, unless it is unavailable in which case the next server in the list is used. You can configure the order of servers in the server group. In the WebUI, use the **up** or **down** arrows to order the servers (the top server is the first server in the list). In the CLI, use the **position** parameter to specify the relative order of servers in the list (the lowest value denotes the first server in the list).

As mentioned previously, the first available server in the list is used for authentication. If the server responds with an authentication failure, there is no further processing for the user or client for which the authentication request failed. You can optionally enable *fail-through* authentication for the server group so that if the first server in the list returns an authentication deny, the controller attempts authentication with the next server in the ordered list. The controller attempts authentication with each server in the list until either there is a successful authentication or the list of servers in the group is exhausted. This feature is useful in environments where there are multiple, independent authentication servers; users may fail authentication on one server but can be authenticated on another server.

Before enabling fail-through authentication, note the following:

- This feature is not supported for 802.1x authentication with a server group that consists of external EAP-compliant RADIUS servers. You can, however, use fail-through authentication when the 802.1x authentication is terminated on the controller (AAA FastConnect).
- Enabling this feature for a large server group list may cause excess processing load on the controller. It is recommended that you use server selection based on domain matching whenever possible (see Configuring Dynamic Server Selection on page 239).
- Certain servers, such as the RSA RADIUS server, lock out the controller if there are multiple authentication failures. Therefore you should not enable fail-through authentication with these servers.

In the following example, you create a server group "corp-serv" with two LDAP servers (ldap-1 and ldap-2), each of which contains a subset of the usernames and passwords used in the network. When you enable fail-through authentication, users that fail authentication on the first server in the server list will be authenticated with the second server.

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **LDAP Server** to display the LDAP Server List.
3. Enter **ldap-1** for the server name and click **Add**.
4. Enter **ldap-2** for the server name and click **Add**.
5. Under the Servers tab, select **ldap-1** to configure server parameters. Enter the IP address for the server. Select the **Mode** checkbox to activate the authentication server. Click **Apply**.
6. Repeat step 5 on page 239 to configure ldap-2.
7. Display the Server Group list: Under the Servers tab, select **Server Group**.
8. Enter **corp-serv** as the new server group and click **Add**.
9. Select **corp-serv,** under the Server tab, to configure the server group.
10. Select **Fail Through**.
11. Under Servers, click **New** to add a server to the group. Select ldap-1 from the drop-down list and click **Add Server**.
12. Repeat step 11 on page 239 to add ldap-2 to the group.
13. Click **Apply**.

Using the CLI

```
(host)(config) #aaa authentication-server ldap ldap-1
  host 10.1.1.234
(host)(config) #aaa authentication-server ldap ldap-2
  host 10.2.2.234

(host)(config) #aaa server-group corp-serv
  auth-server ldap-1 position 1
  auth-server ldap-2 position 2
  allow-fail-through
```

Configuring Dynamic Server Selection

The controller can dynamically select an authentication server from a server group based on the user information sent by the client in an authentication request. For example, an authentication request can include client or user information in one of the following formats:

- <domain>\<user> : for example, corpnet.com\darwin
- <user>@<domain> : for example, darwin@corpnet.com
- host/<pc-name>.<domain> : for example, host/darwin-g.finance.corpnet.com (this format is used with 802.1x machine authentication in Windows environments)

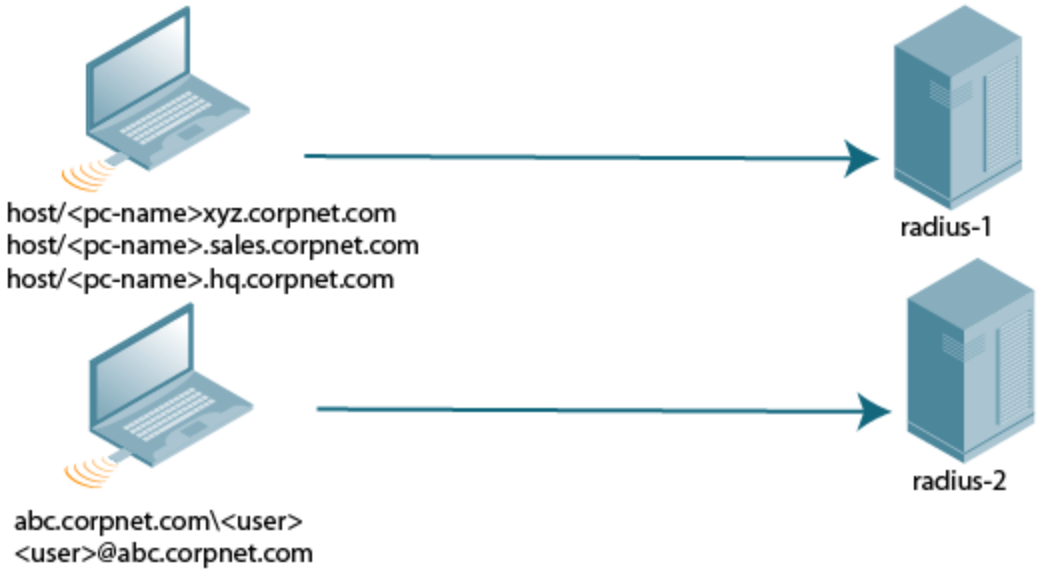
When you configure a server in a server group, you can optionally associate the server with one or more match rules. A match rule for a server can be one of the following:

- The server is selected if the client/user information *contains* a specified string.
- The server is selected if the client/user information *begins* with a specified string.
- The server is selected if the client/user information *exactly* matches a specified string.

You can configure multiple match rules for the same server. The controller compares the client/user information with the match rules configured for each server, starting with the first server in the server group. If a match is found, the controller sends the authentication request to the server with the matching rule. If no match is found before the end of the server list is reached, an error is returned and no authentication request for the client/user is sent.

For example, Figure 29 depicts a network consisting of several subdomains in corpnet.com. The server radius-1 provides 802.1x machine authentication to PC clients in xyz.corpnet.com, sales.corpnet.com, and hq.corpnet.com. The server radius-2 provides authentication for users in abc.corpnet.com.

Figure 29 *Domain-Based Server Selection Example*



You configure the following rules for servers in the corp-serv server group:

- radius-1 is selected if the client information starts with "host."
- radius-2 is selected if the client information contains "abc.corpnet.com."

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Under the Servers tab, select **Server Group** to display the Server Group list.
3. Enter **corp-serv** for the new server group and click **Add**.
4. Under the Servers tab, select **corp-serv** to configure the server group.
5. Under Servers, click **New** to add the radius-1 server to the group. Select radius-1 from the drop-down list.
   a. For Match Type, select **Authstring**.
   b. For Operator, select **starts-with**.
   c. For Match String, enter **host/**.
   d. Click **Add Rule >>**.
   e. Scroll to the right and click **Add Server**.
6. Under Servers, click **New** to add the radius-2 server to the group. Select radius-2 from the drop-down list.
   a. For Match Type, select **Authstring.**
   b. For Operator, select **contains**.
   c. For Match String, enter **abc.corpnet.com.**

d. Click **Add Rule >>**.

e. Scroll to the right and click **Add Server**.



The last server you added to the server group (radius-2) automatically appears as the first server in the list. In this example, the order of servers is not important. If you need to reorder the server list, scroll to the right and click the up or down arrow for the appropriate server.

7. Click **Apply**.

Using the CLI

```
(host)(config) #aaa server-group corp-serv
   auth-server radius-1 match-authstring starts-with host/ position 1
   auth-server radius-2 match-authstring contains abc.corpnet.com position 2
```

Configuring Match FQDN Option

You can also use the "match FQDN" option for a server match rule. With a match FQDN rule, the server is selected if the <domain> portion of the user information in the formats <domain>\<user> or <user>@<domain> *exactly* matches a specified string. Note the following caveats when using a match FQDN rule:

- This rule does *not* support client information in the host/<pc-name>.<domain> format, so it is not useful for 802.1x machine authentication.
- The match FQDN option performs matches on only the <domain> portion of the user information sent in an authentication request. The match-authstring option (described previously) allows you to match all or a portion of the user information sent in an authentication request.

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Under the Servers tab, select **Server Group** to display the Server Group list.
3. Enter **corp-serv** for the new server group and click **Add**.
4. Under the Servers tab, select **corp-serv** to configure the server group.
5. Under Servers, click **New** to add the radius-1 server to the group. Select radius-1 from the drop-down list.
   a. For Match Type, select **FQDN**.
   b. For Match String, enter **corpnet.com**.
   c. Click **Add Rule >>**.
   d. Scroll to the right and click **Add Server**.
6. Click **Apply**.

Using the CLI

```
(host)(config) #aaa server-group corp-serv
   auth-server radius-1 match-fqdn corpnet.com
```

Trimming Domain Information from Requests

Before the controller forwards an authentication request to a specified server, it can truncate the domain-specific portion of the user information. This is useful when user entries on the authenticating server do not include domain information. You can specify this option with any server match rule. This option is only applicable when the user information is sent to the controller in the following formats:

- <domain>\<user> : the <domain>\ portion is truncated
- <user>@<domain> : the @<domain> portion is truncated

This option does not support client information sent in the format host/<pc-name>.<domain>

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **Server Group** to display the Server Group list.
3. Enter the name of the new server group and click **Add**.
4. Select the name to configure the server group.
5. Under Servers, click **Edit** for a configured server or click **New** to add a server to the group.
   - If editing a configured server, select Trim FQDN, scroll right, and click **Update Server**.
   - If adding a new server, select a server from the drop-down list, then select Trim FQDN, scroll right, and click **Add Server**.
6. Click **Apply**.

Using the CLI

```
(host)(config) #aaa server-group corp-serv
  auth-server radius-2 match-authstring contains abc.corpnet.com trim-fqdn
```

Configuring Server-Derivation Rules

When you configure a server group, you can set the VLAN or role for clients based on attributes returned for the client by the server during authentication. The server derivation rules apply to all servers in the group. The user role or VLAN assigned through server derivation rules takes precedence over the default role and VLAN configured for the authentication method.

The authentication servers must be configured to return the attributes for the clients during authentication. For instructions on configuring the authentication attributes in a Windows environment using IAS, refer to the documentation attechnet2.microsoft.com/windowsserver/en/technologies/ias.mspx

The server rules are applied based on the first match principle. The first rule that is applicable for the server and the attribute returned is applied to the client, and would be the only rule applied from the server rules. These rules are applied uniformly across all servers in the server group.

Table 43 describes the server rule parameters you can configure.

Table 43: *Server Rule Configuration Parameters*

Parameter	Description
Role or VLAN	The server derivation rules can be for either user role or VLAN assignment. With Role assignment, a client can be assigned a specific role based on the attributes returned. In case of VLAN assignment, the client can be placed in a specific VLAN based on the attributes returned.
Attribute	This is the attribute returned by the authentication server that is examined for *Operation* and *Operand* match.
Operation	This is the match method by which the string in *Operand* is matched with the attribute value returned by the authentication server. • contains : The rule is applied if and only if the attribute value contains the string in parameter *Operand.* • starts-with : The rule is applied if and only if the attribute value returned

Parameter	Description
	starts with the string in parameter *Operand.* • ends-with : The rule is applied if and only if the attribute value returned ends with the string in parameter *Operand.* • equals : The rule is applied if and only if the attribute value returned equals the string in parameter *Operand.* • not-equals : The rule is applied if and only if the attribute value returned is not equal to the string in parameter *Operand.* • value-of : This is a special condition. What this implies is that the role or VLAN is set to the value of the attribute returned. For this to be successful, the role and the VLAN ID returned as the value of the attribute selected must be already configured on the controller when the rule is applied.
Operand	This is the string to which the value of the returned attribute is matched.
Value	The user role or the VLAN name applied to the client when the rule is matched.
position	Position of the condition rule. Rules are applied based on the first match principle. One is the top. Default: bottom

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **Server Group** to display the Server Group list.
3. Enter the name of the new server group and click **Add**.
4. Select the name to configure the server group.
5. Under Servers, click **New** to add a server to the group.
   a. Select a server from the drop-down list and click **Add**.
   b. Repeat the above step to add other servers to the group.
6. Under Server Rules, click **New** to add server derivation rules for assigning a user role or VLAN.
   a. Enter the attribute.
   b. Select the operation from the drop-down list.
   c. Enter the operand.
   d. To set the role, select **set role** from the **Set** drop-down list and enter the value to be assigned from the **Value** drop-down list.
   e. Or, to set the vlan, select **set vlan** from the **Set** drop-down list and select the VLAN name or ID from the **Value** drop-down list and click the left-arrow.
   f. Click **Add**.
   g. Repeat the above steps to add other rules for the server group.
7. Click **Apply**.

Using the CLI

```
(host) (config) #aaa server-group name
  (host) (Server Group name) #set {role|vlan} condition condition contains operand set-value
  <
  set-value-str> position number
```

Configuring a Role Derivation Rule for the Internal Database

When you add a user entry in the controller's internal database, you can optionally specify a user role (see Managing the Internal Database on page 235). The role specified in the internal database entry to be assigned

to the authenticated client, you must configure a server derivation rule as shown in the following sections:

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **Server Group** to display the Server Group list.
3. Select the **internal** server group.
4. Under Server Rules, click **New** to add a server derivation rule.
   a. For Condition, enter **Role**.
   b. Select **value-of from** the drop-down list.
   c. Select **Set Role** from the drop-down list.
   d. Click **Add**.
5. Click **Apply**.

Using the CLI

```
(host)(config) #aaa server-group internal
   set role condition Role value-of
```

Assigning Server Groups

You can create server groups for the following purposes:

- user authentication
- management authentication
- accounting

You can configure all types of servers for user and management authentication (see Table 44). Accounting is only supported with RADIUS and TACACS+ servers when RADIUS or TACACS+ is used for authentication.

Table 44: *Server Types and Purposes*

	RADIUS	TACACS+	LDAP	Internal Database
User authentication	Yes	Yes	Yes	Yes
Management authentication	Yes	Yes	Yes	Yes
Accounting	Yes	Yes	No	No

User Authentication

For information about assigning a server group for user authentication, refer to the *Roles and Policies* chapter of the *Dell Networking W-Series ArubaOS User Guide*.

Management Authentication

Users who need to access the controller to monitor, manage, or configure the Dell user-centric network can be authenticated with RADIUS, TACACS+, or LDAP servers or the internal database.



Only user record attributes are returned upon a successful authentication. Therefore, to derive a different management role other than the default mgmt auth role, set the server derivation rule based on the user attributes.

Using the WebUI

1. Navigate to the **Configuration > Management > Administration** page.
2. Under the **Management Authentication Servers** section, select the following:
   - **Enable** checkbox
   - **Server Group**
3. Click **Apply**.

Using the CLI

```
(host)(config) #aaa authentication mgmt
   server-group <group>
   enable
```

Accounting

You can configure accounting for RADIUS and TACACS+ server groups.



RADIUS or TACACS+ accounting is only supported when RADIUS or TACACS+ is used for authentication.

RADIUS Accounting

RADIUS accounting allows user activity and statistics to be reported from the controller to RADIUS servers:

1. The controller generates an Accounting Start packet when a user logs in. The code field of transmitted RADIUS packet is set to 4 (Accounting-Request). Note that sensitive information, such as user passwords, are not sent to the accounting server. The RADIUS server sends an acknowledgement of the packet.
2. The controller sends an Accounting Stop packet when a user logs off; the packet information includes various statistics such as elapsed time, input and output bytes, and packets. The RADIUS server sends an acknowledgment of the packet.

The following is the list of attributes that the controller can send to a RADIUS accounting server:

- Acct-Status-Type: This attribute marks the beginning or end of accounting record for a user. Current values are Start and Stop.
- User-Name: Name of user.
- Acct-Session-Id: A unique identifier to facilitate matching of accounting records for a user. It is derived from the user name, IP address, and MAC address. This is set in all accounting packets.
- Acct-Authentic: This indicates how the user was authenticated. Current values are 1 (RADIUS), 2 (Local) and 3 (LDAP).
- Acct-Session-Time: The elapsed time, in seconds, that the client was logged in to the controller. This is only sent in Accounting-Request records where the Acct-Status-Type is Stop.
- Acct-Terminate-Cause: Indicates how the session was terminated and is sent in Accounting-Request records where the Acct-Status-Type is Stop. Possible values are:

  1: User logged off

  4: Idle Timeout

  5: Session Timeout. Maximum session length timer expired.

  7: Admin Reboot: Administrator is ending service, for example prior to rebooting the controller.
- NAS-Identifier: This is set in the RADIUS server configuration.
- NAS-IP-Address: IP address of the master controller. You can configure a “global” NAS IP address: in the WebUI, navigate to the **Configuration > Security > Authentication > Advanced** page; in the CLI, use the, **`ip radius nas-ip`** command.

- NAS-Port: Physical or virtual port (tunnel) number through which the user traffic is entering the controller.
- NAS-Port-Type: Type of port used in the connection. This is set to one of the following:
  - 5: admin login
  - 15: wired user type
  - 19: wireless user
- Framed-IP-Address: IP address of the user.
- Calling-Station-ID: MAC address of the user.
- Called-station-ID: MAC address of the controller.

The following attributes are sent in Accounting-Request packets when Acct-Status-Type value is Start:

- Acct-Status-Type
- User-Name
- NAS-IP-Address
- NAS-Port
- NAS-Port-Type
- NAS-Identifier
- Framed-IP-Address
- Calling-Station-ID
- Called-station-ID
- Acct-Session-ID
- Acct-Authentic

The following attributes are sent in Accounting-Request packets when Acct-Status-Type value is Stop:

- Acct-Status-Type
- User-Name
- NAS-IP-Address
- NAS-Port
- NAS-Port-Type
- NAS-Identifier
- Framed-IP-Address
- Calling-Station-ID
- Called-station-ID
- Acct-Session-ID
- Acct-Authentic
- Terminate-Cause
- Acct-Session-Time

The following attributes are sent only in Accounting Stop packets (they are not sent in Accounting Start packets):

- Acct-Input-Octets
- Acct-Output-Octets
- Acct-Input-Packets
- Acct-Output-Packets

Remote APs in split-tunnel mode now support RADIUS accounting. If you enable RADIUS accounting in a split-tunnel Remote AP's AAA profile, the controller sends a RADIUS accounting start record to the RADIUS server when a user associates with the remote AP, and sends a stop record when the user logs out or is deleted from the user database. If interim accounting is enabled, the controller sends updates at regular intervals. Each interim record includes cumulative user statistics, including received bytes and packets counters.

You can use either the WebUI or CLI to assign a server group for RADIUS accounting.

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > AAA Profiles** page.
2. Select AAA Profile, then the AAA profile instance.
3. (Optional) In the **Profile Details** pane, select RADIUS Interim Accounting to allow the controller to send Interim-Update messages with current user statistics to the server at regular intervals. This option is disabled by default, allowing the controller to send only *start* and *stop* messages RADIUS accounting server.
4. In the profile list, scroll down and select the Radius Accounting Server Group for the AAA profile. Select the server group from the drop-down list.

   You can add additional servers to the group or configure server rules.
5. Click **Apply**.

Using the CLI

```
(host)(config) #aaa profile <profile>
   radius-accounting <group>
   radius-interim-accounting
```

RADIUS Accounting on Multiple Servers

ArubaOS provides support for the controllers to send RADIUS accounting to multiple RADIUS servers. The controller notifies all the RADIUS servers to track the status of authenticated users. Accounting messages are sent to all the servers configured in the server group in a sequential order.

You can enable multiple server account functionality by using CLI and WebUI:

Using the CLI:

To enable RADIUS Accounting on Multiple Servers functionality, use the following CLI:

```
(host) (config) # aaa profile <profile_name>
   multiple-server-accounting
```

Using the WebUI:

1. Navigate to the **Configuration > Security > Authentication > AAA Profiles** page.
2. Select **AAA Profile**, then select the AAA profile instance.
3. Select **Multiple Server Accounting** checkbox.
4. Click **Apply**.

TACACS+ Accounting

TACACS+ accounting allows commands issued on the controller to be reported to TACACS+ servers. You can specify the types of commands that are reported (action, configuration, or show commands) or have all commands reported.

You can configure TACACS+ accounting only with the CLI:

```
(host)(config) #aaa tacacs-accounting server-group <group> command
{action|all|configuration|show} mode {enable|disable}
```

Configuring Authentication Timers

Table 45 describes the timers you can configure that apply to all clients and servers. These timers can be left at their default values for most implementations.

Table 45: *Authentication Timers*

Timer	Description
User Idle Timeout	Maximum period after which a client is considered idle if there is no wireless traffic from the client.The timeout period is reset if there is wireless traffic. If there is no wireless traffic in the timeout period, the client is aged out. Once the timeout period has expired, the user is removed. If the keyword **seconds** is not specified, the value defaults to minutes at the command line. Range: 1–255 minutes (30–15300 seconds) Default: 5 minutes (300 seconds)
Authentication Server Dead Time	Maximum period, in minutes, that the controller considers an unresponsive authentication server to be "out of service." This timer is only applicable if there are two or more authentication servers configured on the controller. If there is only one authentication server configured, the server is never considered out of service and all requests are sent to the server. If one or more backup servers are configured and a server is unresponsive, it is marked as out of service for the dead time; subsequent requests are sent to the next server on the priority list for the duration of the dead time. If the server is responsive after the dead time has elapsed, it can take over servicing requests from a lower-priority server; if the server continues to be unresponsive, it is marked as down for the dead time. Range: 0–50 minutes Default: 10 minutes
Logon User Lifetime	Maximum time, in minutes, unauthenticated clients are allowed to remain logged on. Range: 0–255 minutes Default: 5 minutes
User Interim stats frequency	Set the timeout value for user stats reporting in minutes or seconds. Range ; 300-600 seconds, or 5-10 minutes Default : 600 seconds

Setting an Authentication Timer

To set an authentication timer, complete one of the following procedures:

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Advanced** page.
2. Configure the timers as described above.
3. Click **Apply** before moving on to another page or closing the browser window. If you do not perform this step, you will lose your configuration changes.

Using the CLI

The commands below configure timers you can apply to clients. If the optional seconds keyword is not specified for the **idle-timeout** and **stats-timeout** parameters, the value defaults to minutes.

```
(host)(config) #aaa timers
  dead-time <minutes>
```

```
   idle-timeout <time> [seconds]
   logon-lifetime <0-255>
   stats-timeout <time> [seconds]
```

Authentication Server Load Balancing

Load balancing of authentication servers ensures that the authentication load is split across multiple authentication servers, thus avoiding any one particular authentication server from being overloaded. Authentication Server Load Balancing functionality enables the Dell Mobility Controller to perform load balancing of authentication requests destined to external authentication servers (Radius/LDAP etc). This prevents any one authentication server from having to handle the full load during heavy authentication periods, such as at the start of the business day.

Earlier, the controller used the first authentication server in the server group list. The remaining servers in that server group will be used in sequential order only when an authentication server is down. Thus, the controllers performed fail-over and not load balancing of authentication servers.

The load balancing algorithm computes the expected time taken to authenticate a new client for each authentication server and chooses that authentication server for which the expected authentication time is least. The load balancing algorithm maintains re-authentication stickiness i.e. at the time of re-authentication the request will be forwarded to the same server where it was authenticated earlier.

Enabling Authentication Server Load Balancing Functionality

A new **load–balancing enable** parameter has been introduced in **aaa server-group test** command to enable authentication server load balancing functionality.

```
aaa server-group <sg_name>
   load-balance
   auth-server s1
   auth-server s2
```

You can use the following command to disable load balancing:

```
aaa server-group<sg_name>
   no load-balance
```



If you configure internal server in the server group, load balancing is not applicable to the internal server. Internal server will be used as fall back when all the other servers in the server group are down.

This chapter describes how to configure MAC-based authentication on the Dell controller using the WebUI.

Use MAC-based authentication to authenticate devices based on their physical media access control (MAC) address. While not the most secure and scalable method, MAC-based authentication implicitly provides an addition layer of security authentication devices. MAC-based authentication is often used to authenticate and allow network access through certain devices while denying access to the rest. For example, if clients are allowed access to the network via station A, then one method of authenticating station A is MAC-based. Clients may be required to authenticate themselves using other methods depending on the network privileges required.

MAC-based authentication can also be used to authenticate Wi-Fi phones as an additional layer of security to prevent other devices from accessing the voice network using what is normally an insecure SSID.

This chapter describes the following topics:

- Configuring MAC-Based Authentication on page 250
- Configuring Clients on page 251

Configuring MAC-Based Authentication

Before configuring MAC-based authentication, you must configure the following options:

- User role—The user role that will be assigned as the default role for the MAC-based authenticated clients. (See Roles and Policies on page 364 for information on firewall policies to configure roles.)

  Configure the default user role for MAC-based authentication in the AAA profile. If derivation rules exist or if the client configuration in the internal database has a role assignment, these values take precedence over the default user role.
- Authentication server group— The authentication server group that the controller uses to validate the clients. The internal database can be used to configure the clients for MAC-based authentication. See Configuring Clients on page 251 for information on configuring the clients on the local database. For information on configuring authentication servers and server groups, see Authentication Servers on page 225.

Configuring the MAC Authentication Profile

Table 46 describes the parameters you can configure for MAC-based authentication.

Table 46: *MAC Authentication Profile Configuration Parameters*

Parameter	Description
Delimiter	Delimiter used in the MAC string: • colon specifies the format Xx:XX:XX:XX:XX:XX • dash specifies the format XX-XX-XX-XX-XX-XX • none specifies the format XXXXXXXXXXXX • oui-nic specifies the format XXXXXX:XXXXXX Default: none **NOTE:** This parameter is available for the **aaa authentication-server radius** command.
Case	The case (upper or lower) used in the MAC string. Default: lower
Max Authentication failures	Number of times a station can fail to authenticate before it is blacklisted. A value of zero disables blacklisting. Default: zero (0)

Using the WebUI to configure a MAC authentication profile

1. Navigate to the **Configuration > Security > Authentication > L2 Authentication** page.
2. Select MAC Authentication Profile.
3. Enter a profile name and click **Add**.
4. Select the profile name to display configurable parameters.
5. Configure the parameters, as described in Table 46.
6. Click **Apply**.

Using the CLI to configure a MAC authentication profile

Execute the following command to configure a MAC authentication profile:

```
(host)(configure) #aaa authentication mac <profile>
   case {lower|upper}
   delimiter {colon|dash|none}
   max-authentication-failures <number>
```

Configuring Clients

You can create entries in the controller's internal database to authenticate client MAC addresses. The internal database contains a list of clients along with the password and default role for each client. To configure entries in the internal database for MAC authentication, you enter the MAC address for both the username and password for each client.



You must enter the MAC address using the delimiter format configured in the MAC authentication profile. The default delimiter is none, which means that MAC addresses should be in the format xxxxxxxxxxxx. If you specify colons for the delimiter, you can enter MAC addresses in the format xx:xx:xx:xx:xx:xx.

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **Internal DB**.
3. Click **Add User** in the **Users** section. The user configuration page displays.

4. For **User Name** and **Password**, enter the MAC address for the client. Use the format specified by the Delimiter parameter in the MAC Authentication profile. For example, if the MAC Authentication profile specifies the default delimiter (none), enter MAC addresses in the format `xxxxxxxxxxxx`.
5. Click **Enabled** to activate this entry on creation.
6. Click **Apply**.



The configuration does not take effect until you perform this step.

In the CLI

Enter the following command in enable mode:

```
(host)(config) #local-userdb add username <macaddr> password <macaddr>...
```

802.1X is an Institute of Electrical and Electronics Engineers (IEEE) standard that provides an authentication framework for WLANs. 802.1x uses the Extensible Authentication Protocol (EAP) to exchange messages during the authentication process. The authentication protocols that operate inside the 802.1X framework that are suitable for wireless networks include EAP-Transport Layer Security (EAP-TLS), Protected EAP (PEAP), and EAP-Tunneled TLS (EAP-TTLS). These protocols allow the network to authenticate the client while also allowing the client to authenticate the network.

This chapter describes the following topics:

- Understanding 802.1X Authentication on page 253
- Configuring 802.1X Authentication on page 256
- Sample Configurations on page 265
- Performing Advanced Configuration Options for 802.1X on page 281

Other types of authentication not discussed in this section can be found in the following sections of this guide:

- Captive portal authentication: Configuring Captive Portal Authentication Profiles on page 312
- VPN authentication: Planning a VPN Configuration on page 337
- MAC authentication: Configuring MAC-Based Authentication on page 250
- Stateful 802.1x, stateful NTLM, and WISPr authentication: Stateful and WISPr Authentication on page 285

Understanding 802.1X Authentication

802.1x authentication consists of three components:

- The *supplicant*, or client, is the device attempting to gain access to the network. You can configure the Dell user-centric network to support 802.1x authentication for wired usersa wireless users.
- The *authenticator* is the gatekeeper to the network and permits or denies access to the supplicants.
- The *Dell controller* acts as the authenticator, relaying information between the authentication server and supplicant. The EAP type must be consistent between the authentication server and supplicant, and is transparent to the controller.

  The authentication server provides a database of information required for authentication, and informs the authenticator to deny or permit access to the supplicant.

  The 802.1X authentication server is typically an EAP-compliant Remote Access Dial-In User Service (RADIUS) server which can authenticate either users (through passwords or certificates) or the client computer.

  An example of an 802.1X authentication server is the Internet Authentication Service (IAS) in Windows (see technet.microsoft.com/en-us/library/cc759077(WS.10).aspx).

  In Dell user-centric networks, you can terminate the 802.1x authentication on the controller. The controller passes user authentication to its internal database or to a “backend” non-802.1X server. This feature, also called *AAA FastConnect*, is useful for deployments where an 802.1X EAP-compliant RADIUS server is not available or required for authentication.

Supported EAP Types

Following is the list of supported EAP types:

- PEAP—Protected EAP (PEAP) is an 802.1X authentication method that uses server-side public key certificates to authenticate clients with the server. The PEAP authentication creates an encrypted SSL / TLS tunnel between the client and the authentication server. The exchange of information is encrypted and stored in the tunnel to ensure that the user credentials are kept secure.
- EAP-GTC—The EAP-GTC (Generic Token Card) type uses clear text method to exchange authentication controls between the client and the server. Since the authentication mechanism uses the one-time tokens (generated by the card), this method of credential exchange is considered safe. In addition, EAP-GTC is used in PEAP or TTLS tunnels in wireless environments. The EAP-GTC is described in RFC 2284.
- EAP-AKA—The EAP-AKA (Authentication and Key Agreement) authentication mechanism is typically used in mobile networks that include Universal Mobile Telecommunication Systems (UMTS) and CDMA 2000. This method uses the information stored in the Subscriber Identity Module (SIM) for authentication. The EAP-AKA is described in RFC 4187.
- EAP-FAST—The EAP-FAST (Flexible Authentication via Secure Tunneling) is an alternative authentication method to PEAP. This method uses the Protected Access Credential (PAC) for verifying clients on the network. The EAP-FAST is described in RFC 4851.
- EAP-MD5—The EAP-MD5 method verifies MD5 hash of a user password for authentication. This method is commonly used in a trusted network. The EAP-MD5 is described in RFC 2284.
- EAP-POTP—The EAP type 32 is supported. Complete details are described in RFC 4793.
- EAP-SIM—The EAP-SIM (Subscriber Identity Module) uses Global System for Mobile Communication (GSM) Subscriber Identity Module (SIM) for authentication and session key distribution. This authentication mechanism includes network authentication, user anonymity support, result indication, and fast re-authentication procedure. Complete details about this authentication mechanism is described in RFC 4186.
- EAP-TLS—The EAP-TLS (Transport Layer Security) uses Public key Infrastructure (PKI) to set up authentication with a RADIUS server or any authentication server. This method requires the use of a client-side certificate for communicating with the authentication server. The EAP-TLS is described in RFC 5216.
- EAP-TLV—The EAP-TLV (type-length-value) method allows you to add additional information in an EAP message. Often this method is used to provide more information about an EAP message such as status information or authorization data. This method is always used after a typical EAP authentication process.
- EAP-TTLS—The EAP-TTLS (Tunneled Transport Layer Security) method uses server-side certificates to set up authentication between clients and servers. The actual authentication is, however, performed using passwords. Complete details about EAP-TTLS is described in RFC 5281.
- LEAP—Lightweight Extensible Authentication Protocol (LEAP) uses dynamic WEP keys and mutual authentication between the client and the RADIUS server.
- ZLXEAP—ZoneLabs EAP is an EAP method that has been allocated EAP Type 44 by IANA. For more information, visit tools.ietf.org/html/draft-bersani-eap-synthesis-sharedkeymethods-00#page-30.

Configuring Authentication with a RADIUS Server

See Table 47 for an overview of the parameters that you need to configure on authentication components when the authentication server is an 802.1X EAP-compliant RADIUS server.

Figure 30 *802.1X Authentication with RADIUS Server*



The supplicant and the authentication server must be configured to use the same EAP type. The controller does not need to know the EAP type used between the supplicant and authentication server.

For the controller to communicate with the authentication server, you must configure the IP address, authentication port, and accounting port of the server on the controller. The authentication server must be configured with the IP address of the RADIUS client, which is the controller in this case. Both the controller and the authentication server must be configured to use the same shared secret.

NOTE

Additional information on EAP types supported in a Windows environment, Microsoft supplicants, and authentication servers, is available at technet.microsoft.com/en-us/library/cc782851(WS.10).aspx.

The client communicates with the controller through a GRE tunnel to form an association with an AP and to get authenticated in the network. Therefore, the network authentication and encryption configured for an ESSID must be the same on both the client and the controller.

Configuring Authentication Terminated on Controller

User authentication is performed either via the controller's internal database or a non-802.1X server. See 802.1x Authentication Profile Basic WebUI Parameters on page 257 for an overview of the parameters that you need to configure on 802.1X authentication components when 802.1X authentication is terminated on the controller (AAA FastConnect).

Figure 31 *802.1X Authentication with Termination on Controller*



In this scenario, the supplicant is configured for EAP-Transport Layer Security (TLS) or EAP-Protected EAP (PEAP).

- EAP-TLS is used with smart card user authentication. A smart card holds a digital certificate which, with the user-entered personal identification number (PIN), allows the user to be authenticated on the network. EAP-TLS relies on digital certificates to verify the identities of both the client and the server.

  EAP-TLS requires that you import server and certification authority (CA) certificates onto the controller (see Configuring and Using Certificates with AAA FastConnect on page 261). The client certificate is verified on the controller (the client certificate must be signed by a known CA) before the username is checked on the authentication server.
- EAP-PEAP uses TLS to create an encrypted tunnel. Within the tunnel, one of the following "inner EAP" methods is used:
  - EAP-Generic Token Card (GTC): Described in RFC 2284, this EAP method permits the transfer of unencrypted usernames and passwords from client to server. The main uses for EAP-GTC are one-time token cards such as SecureID and the use of an LDAP or RADIUS server as the user authentication server. You can also enable caching of user credentials on the controller as a backup to an external authentication server.
  - EAP-Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2): Described in RFC 2759, this EAP method is widely supported by Microsoft clients. A RADIUS server must be used as the backend authentication server.

If you use the controller's internal database for user authentication, you need to add the names and passwords of the users to be authenticated. If you use an LDAP server for user authentication, you need to configure both the LDAP server and the user IDs and passwords on the controller. If you use a RADIUS server for user authentication, you need to configure the RADIUS server on the controller.

Configuring 802.1X Authentication

On the controller, use the following steps to configure a wireless network that uses 802.1x authentication:

1. Configure the VLANs to which the authenticated users will be assigned. See Network Configuration Parameters on page 148.
2. Configure policies and roles. You can specify a default role for users who are successfully authenticated using 802.1X. You can also configure server derivation rules to assign a user role based on attributes returned by the authentication server; server-derived user roles take precedence over default roles. For more information about policies and roles, see Roles and Policies on page 364.



The Policy Enforcement Firewall Virtual Private Network (PEFV) module provides identity-based security for wired and wireless users and must be installed on the controller. The stateful firewall allows user classification based on user identity, device type, location, and time of day to provide differentiated access for different classes of users. For information about obtaining and installing licenses, see Software Licenses on page 130.

3. Configure the authentication server(s) and server group. The server can be an 802.1X RADIUS server or, if you use AAA FastConnect, a non-802.1X server or the controller's internal database. If you use EAP-GTC within a PEAP tunnel, you can configure an LDAP or RADIUS server as the authentication server (see Authentication Servers on page 225). If you use EAP-TLS, you need to import server and CA certificates on the controller (see Configuring and Using Certificates with AAA FastConnect on page 261).
4. Configure the AAA profile:
   - Select the 802.1X default user role.
   - Select the server group you previously configured for the 802.1x authentication server group.
5. Configure the 802.1X authentication profile. See In the WebUI on page 276.
6. Configure the virtual AP profile for an AP group or for a specific AP:
   - Select the AAA profile you previously configured.
   - In the SSID profile, configure the WLAN for 802.1X authentication.

For details on how to complete the above steps, see Sample Configurations on page 265.

In the WebUI

This section describes how to create and configure a new instance of an 802.1X authentication profile in the WebUI or the CLI.

1. Navigate to the **Configuration > Security > Authentication > L2 Authentication** page.
2. In the **Profiles** list, select **802.1X Authentication Profile**.
3. Enter a name for the profile, then click **Add.**
4. Click **Apply**.
5. In the **Profiles** list, select the 802.1X authentication profile you just created.
6. Change the settings described in Table 47 as desired, then click **Apply**.

   The 802.1X authentication profile configuration settings are divided into two tabs—**Basic** and **Advanced**. The **Basic** tab displays only those configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. If you change a setting on one tab, then click and display the other tab without saving your configuration, that setting will revert to its previous value.

Table 47: *802.1x Authentication Profile Basic WebUI Parameters*

Parameter	Description
Basic 802.1x Authentication Settings	
Max authentication failures	Number of times a user can try to log in with wrong credentials after which the user is blacklisted as a security threat. Set to 0 to disable blacklisting, otherwise enter a non-zero integer to blacklist the user after the specified number of failures. Range: 0-5 failures. Default: 0 failure. **NOTE:** This option may require a license.
Enforce Machine Authentication	Select the Enforce Machine Authentication option to require machine authentication. This option is also available on the **Basic** settings tab. **NOTE:** This option may require a license.
Machine Authentication: Default Machine Role	Default role assigned to the user after completing only machine authentication. The default role for this setting is the “guest” role.
Machine Authentication: Default User Role	Default role assigned to the user after 802.1x authentication. The default role for this setting is the “guest” role.
Reauthentication	Select the Reauthentication checkbox to force the client to do a 802.1X reauthentication after the expiration of the default timer for reauthentication. (The default value of the timer is 24 hours.) If the user fails to reauthenticate with valid credentials, the state of the user is cleared. If derivation rules are used to classify 802.1x-authenticated users, then the reauthentication timer per role overrides this setting. This option is disabled by default.
Termination	Select the **Termination** checkbox to allow 802.1X authentication to terminate on the controller. This option is disabled by default.
Termination EAP-Type	If you enable termination, click either EAP-PEAP or EAP-TLS to select a Extensible Authentication Protocol (EAP) method.
Termination Inner EAP-Type	If you use EAP-PEAP as the EAP method, specify one of the following inner EAP types: • **eap-gtc**: Described in RFC 2284, this EAP method permits the transfer of unencrypted usernames and passwords from client to server. The main uses for EAP-GTC are one-time token cards such as SecureID and the use of LDAP or RADIUS as the user authentication server. You can also enable caching of user credentials on the controller as a backup to an external authentication server. • **eap-mschapv2**: Described in RFC 2759, this EAP method is widely supported by Microsoft clients.
Enforce Suite-B 128 bit or more security level Authentication	Configure Suite-B 128 bit or more security level authentication enforcement.
Enforce Suite-B 128 bit or more security level Authentication	Configure Suite-B 192 bit security level authentication enforcement.
Advanced 802.1x Authentication Settings	

Table 47: *802.1x Authentication Profile Basic WebUI Parameters*

Parameter	Description
Machine Authentication Cache Timeout	The timeout, in hours, for machine authentication. The allowed range of values is 1-1000 hours, and the default value is 24 hours.
Blacklist on Machine Authentication Failure	Select the **Blacklist on Machine Authentication Failure** checkbox to blacklist a client if machine authentication fails. This setting is disabled by default.
Interval between Identity Requests	Interval, in seconds, between identity request retries. Range: 1-65535 seconds. Default: 30 seconds.
Quiet Period after Failed Authentication	The enforced quiet period interval, in seconds, following failed authentication. Range: 1-65535 seconds. Default: 30 seconds.
Reauthentication Interval	Interval, in seconds, between reauthentication attempts. Range: 60-864000 seconds. Default: 86400 seconds (1 day).
Use Server provided Reauthentication Interval	Select this option to override any user-defined reauthentication interval and use the reauthentication period defined by the authentication server.
Multicast Key Rotation Time Interval	Interval, in seconds, between multicast key rotation. Range: 60-864000 seconds. Default: 1800 seconds.
Unicast Key Rotation Time Interval	Interval, in seconds, between unicast key rotation. Range: 60-864000 seconds. Default: 900 seconds.
Authentication Server Retry Interval	Server group retry interval, in seconds. Range: 5-65535 seconds. Default: 30 seconds.
Authentication Server Retry Count	Maximum number of authentication requests that are sent to server group. Range: 0-3 requests. Default: 2 requests.
Framed MTU	Sets the framed Maximum Transmission Unit (MTU) attribute sent to the authentication server. Range: 500-1500 bytes. Default: 1100 bytes.
Number of times ID-Requests are retried	Maximum number of times ID requests are sent to the client. Range: 1-10 retries. Default: 3 retries.
Maximum Number of Reauthentication Attempts	Number of times a user can try to log in with wrong credentials after which the user is blacklisted as a security threat. Set to 0 to disable blacklisting, otherwise enter a value from 0-5 to blacklist the user after the specified number of failures. **NOTE:** If changed from its default value, this option may require a license.
Maximum number of times Held State can be bypassed	Number of consecutive authentication failures which, when reached, causes the controller to not respond to authentication requests from a client while the controller is in a held state after the authentication failure. Before this number is reached, the controller responds to authentication requests from the client even while the controller is in its held state.

Table 47: *802.1x Authentication Profile Basic WebUI Parameters*

Parameter	Description
	(This parameter is applicable when 802.1X authentication is terminated on the controller, also known as AAA FastConnect.) The allowed range of values for this parameter is 0-3 failures, and the default value is 0.
Dynamic WEP Key Message Retry Count	Set the Number of times WPA/WPA2 Key Messages are retried. Range: 1-5 retries. Default: 3 retries.
Dynamic WEP Key Size	The default dynamic WEP key size is 128 bits, If desired, you can change this parameter to 40 bits.
Interval between WPA/WPA2 Key Messages	Interval, in milliseconds, between each WPA key exchanges. Range: 1000-5000 ms. Default: 1000 ms.
Delay between EAP-Success and WPA2 Unicast Key Exchange	Interval, in milliseconds, between EAP-Success and unicast key exchanges. Range: 0-2000 ms. Default: 0 ms (no delay).
Delay between WPA/WPA2 Unicast Key and Group Key Exchange	Interval, in milliseconds, between unicast and multicast key exchange. Time interval in milliseconds. Range: 0-2000. Default: 0 (no delay).
Time interval after which the PMKSA will be deleted	The time interval after which the PMKSA (Pairwise Master Key Security Association) cache is deleted. Time interval in Hours. Range: 1-2000. Default: 8.
WPA/WPA2 Key Message Retry Count	Number of times WPA/WPA2 key messages are retried. Range: 1-5 retries. Default: 3 retries.
Multicast Key Rotation	Select this checkbox to enable multicast key rotation. This feature is disabled by default.
Unicast Key Rotation	Select this checkbox to enable unicast key rotation. This feature is disabled by default.
Opportunistic Key Caching	By default, the 802.1X authentication profile enables a cached pairwise master key (PMK) which is derived through a client and an associated AP. This key is used when the client roams to a new AP. This allows clients faster roaming without a full 802.1x authentication. Uncheck this option to disable this feature. **NOTE:** Make sure that the wireless client (the 802.1X supplicant) supports this feature. If the client does not support this feature, the client will attempt to renegotiate the key whenever it roams to a new AP. As a result, the key cached on the controller can be out of sync with the client's key.
Validate PMKID	This parameter instructs the controller to check the pairwise master key (PMK) ID sent by the client. When you enable this option, the client must send a PMKID in the associate or reassociate frame to indicate that it supports OKC or PMK caching; otherwise, full 802.1x authentication takes place. **NOTE:** This feature is optional, since most clients that support OKC and PMK caching do not send the PMKID in their association request.

Table 47: *802.1x Authentication Profile Basic WebUI Parameters*

Parameter	Description
Use Session Key	Select the **Use Session Key** option to use the RADIUS session key as the unicast WEP key. This option is disabled by default.
Use Static Key	Select the **Use Static Key** option to use a static key as the unicast/multicast WEP key. This option is disabled by default.
xSec MTU	Set the maximum transmission unit (MTU) for frames using the xSec protocol. Range: 1024-1500 bytes. Default: 1300 bytes.
Token Caching	If you select EAP-GTC as the inner EAP method, you can select the **Token Caching** checkbox to enable the controller to cache the username and password of each authenticated user. The controller continues to reauthenticate users with the remote authentication server. However, if the authentication server is unavailable, the controller will inspect its cached credentials to reauthenticate users. This option is disabled by default.
Token Caching Period	If you select EAP-GTC as the inner EAP method, you can specify the timeout period, in hours, for the cached information. The default value is 24 hours.
CA-Certificate	Click the **CA-Certificate** drop-down list and select a certificate for client authentication. The CA certificate needs to be loaded in the controller before it will appear on this list.
Server-Certificate	Click the **Server-Certificate** drop-down list and select a server certificate the controller will use to authenticate itself to the client.
TLS Guest Access	Select **TLS Guest Access** to enable guest access for EAP-TLS users with valid certificates. This option is disabled by default.
TLS Guest Role	Click the **TLS Guest Role** drop-down list and select the default user role for EAP-TLS guest users. This option may require a license.
Ignore EAPOL-START after authentication	Select **Ignore EAPOL-START** after authentication to ignore EAPOL-START messages after authentication. This option is disabled by default.
Handle EAPOL-Logoff	Select **Handle EAPOL-Logoff** to enable handling of EAPOL-LOGOFF messages. This option is disabled by default.
Ignore EAP ID during negotiation	Select **Ignore EAP ID during negotiation** to ignore EAP IDs during negotiation. This option is disabled by default.
WPA-Fast-Handover	Select this option to enable WPA-fast-handover on phones that support this feature. WAP fast-handover is disabled by default.
Disable rekey and reauthentication for clients on call	This feature disables rekey and reauthentication for VoWLAN clients. It is disabled by default, meaning that rekey and reauthentication is enabled. **NOTE:** This option may require a license This option may require a license.
Check certificate common name against AAA server	If you use client certificates for user authentication, enable this option to verify that the certificate's common name exists in the server. This parameter is enabled by default in the default-cap and default-rap VPN profiles, and disabled by default on all other VPN profiles.

In the CLI

The following command configures settings for an 802.1X authentication profiles. Individual parameters are described in the previous table.

```
(host)(config) #aaa authentication dot1x {<profile>|countermeasures}
   ca-cert <certificate>
   clear
   clone <profile>
   eapol-logoff
   framed-mtu <mtu>
   heldstate-bypass-counter <number>
   ignore-eap-id-match
   ignore-eapolstart-afterauthentication
   machine-authentication blacklist-on-failure|{cache-timeout <hours>}|enable|
   {machine-default-role <role>}|{user-default-role <role>}
   max-authentication-failures <number>
   max-requests <number>
   multicast-keyrotation
   no ...
   opp-key-caching
   reauth-max <number>
   reauthentication
   server {server-retry <number>|server-retry-period <seconds>}
   server-cert <certificate>
   termination {eap-type <type>}|enable|enable-token-caching|{inner-eap-type (eapgtc|
   eap-mschapv2)}|{token-caching-period <hours>}
   timer {idrequest_period <seconds>}|{mkey-rotation-period <seconds>}|{quiet-period
   <seconds>}|{reauth-period <seconds>}|{ukey-rotation-period <seconds>}|{wpagroupkey-
   delay <seconds>}|{wpa-key-period <milliseconds>}
   tls-guest-access
   tls-guest-role <role>
   unicast-keyrotation
   use-session-key
   use-static-key
   validate-pmkid
   voice-aware
   wep-key-retries <number>
   wep-key-size {40|128}
   wpa-fast-handover
   wpa-key-retries <number>
   xSec-mtu <mtu>
```

Configuring and Using Certificates with AAA FastConnect

The controller supports 802.1x authentication using digital certificates for AAA FastConnect.

- Server Certificate—A server certificate installed in the controller verifies the authenticity of the controller for 802.1x authentication. Dell controllers ship with a demonstration digital certificate. Until you install a customer-specific server certificate in the controller, this demonstration certificate is used by default for all secure HTTP connections (such as the WebUI and captive portal) and AAA FastConnect. This certificate is included primarily for the purposes of feature demonstration and convenience, and is not intended for long-term use in production networks. Users in a production environment are urged to obtain and install a certificate issued for their site or domain by a well-known certificate authority (CA). You can generate a Certificate Signing Request (CSR) on the controller to submit to a CA. For information on how to generate a CSR and how to import the CA-signed certificate into the controller, see Managing Certificates on page 795.
- Client Certificates—Client certificates are verified on the controller (the client certificate must be signed by a known CA) before the username is checked on the authentication server. To use client certificate

authentication for AAA FastConnect, you need to import the following certificates into the controller (see Importing Certificates on page 798):

- Controller's server certificate
- CA certificate for the CA that signed the client certificates

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > L2 Authentication** page.
2. In the **Profiles** list, select **802.1x Authentication Profile**.
3. Select the **default** 802.1x authentication profile from the drop-down list to display configuration parameters.
4. In the **Basic** tab, select **Termination**.
5. Select the **Advanced** Tab.
6. In the Server-Certificate field, select the server certificate imported into the controller.
7. In the CA-Certificate field, select the CA certificate imported into the controller.
8. Click **Save As**. Enter a name for the 802.1x authentication profile.
9. Click **Apply**.

In the CLI

Use the following commands to configure the 802.1x profile:

```
(host)(config) #aaa authentication dot1x <profile>
      termination enable
      server-cert <certificate>
      ca-cert <certificate>
```

Configuring User and Machine Authentication

When a Windows device boots, it logs onto the network domain using a machine account. Within the domain, the device is authenticated before computer group policies and software settings can be executed; this process is known as *machine authentication*. Machine authentication ensures that only authorized devices are allowed on the network.

You can configure 802.1x for both user and machine authentication (select the **Enforce Machine Authentication** option described in Table 47). This tightens the authentication process further, since both the device and user need to be authenticated.

Working with Role Assignment with Machine Authentication Enabled

When you enable machine authentication, there are two additional roles you can define in the 802.1x authentication profile:

- Machine authentication default machine role
- Machine authentication default user role

While you can select the same role for both options, you should define the roles as per the polices that need to be enforced. Also, these roles can be different from the 802.1x authentication default role configured in the AAA profile.

With machine authentication enabled, the assigned role depends upon the success or failure of the machine and user authentications. In certain cases, the role that is ultimately assigned to a client can also depend upon attributes returned by the authentication server or server derivation rules configured on the controller.

Table 48 describes role assignment based on the results of the machine and user authentications.

Table 48: *Role Assignment for User and Machine Authentication*

Machine Auth Status	User Auth Status	Description	Role Assigned
Failed	Failed	Both machine authentication and user authentication failed. L2 authentication failed.	No role assigned. No access to the network allowed.
Failed	Passed	Machine authentication failed (for example, the machine information is not present on the server) and user authentication succeeded. Server-derived roles do not apply.	Machine authentication default user role configured in the 802.1x authentication profile.
Passed	Failed	Machine authentication succeeded and user authentication has not been initiated. Server-derived roles do not apply.	Machine authentication default machine role configured in the 802.1x authentication profile.
Passed	Passed	Both machine and user are successfully authenticated. If there are server-derived roles, the role assigned via the derivation take precedence. This is the only case where server-derived roles are applied.	A role derived from the authentication server takes precedence. Otherwise, the 802.1x authentication default role configured in the AAA profile is assigned.

For example, if the following roles are configured:

- 802.1x authentication default role (in AAA profile): dot1x_user
- Machine authentication default machine role (in 802.1x authentication profile): dot1x_mc
- Machine authentication default user role (in 802.1x authentication profile): guest

Role assignment is as follows:

- If both machine and user authentication succeed, the role is dot1x_user. If there is a server-derived role, the server-derived role takes precedence.
- If only machine authentication succeeds, the role is dot1x_mc.
- If only user authentication succeeds, the role is guest.
- On failure of both machine and user authentication, the user does not have access to the network.

With machine authentication enabled, the VLAN to which a client is assigned (and from which the client obtains its IP address) depends upon the success or failure of the machine and user authentications. The VLAN that is ultimately assigned to a client can also depend upon attributes returned by the authentication server or server derivation rules configured on the controller (see Understanding VLAN Assignments on page 157). If machine authentication is successful, the client is assigned the VLAN configured in the virtual AP profile. However, the client can be assigned a derived VLAN upon successful user authentication.



You can optionally assign a VLAN as part of a user role configuration. You should not use VLAN derivation if you configure user roles with VLAN assignments.

Table 49 describes VLAN assignment based on the results of the machine and user authentications when VLAN derivation is used.

Table 49: *VLAN Assignment for User and Machine Authentication*

Machine Auth Status	User Auth Status	Description	VLAN Assigned
Failed	Failed	Both machine authentication and user authentication failed. L2 authentication failed.	No VLAN.
Failed	Passed	Machine authentication failed (for example, the machine information is not present on the server) and user authentication succeeded.	VLAN configured in the virtual AP profile.
Passed	Failed	Machine authentication succeeded and user authentication has not been initiated.	VLAN configured in the virtual AP profile.
Passed	Passed	Both machine and user are successfully authenticated.	Derived VLAN. Otherwise, VLAN configured in the virtual AP profile.



The administrator can now associate a VLAN ID to a client data based on the authentication credentials in a bridge mode.

Enabling 802.1x Supplicant Support on an AP

This release of ArubaOS provides 802.1X supplicant support on the Access Point (AP). The AP can be used as a 802.1x supplicant where access to the wired Ethernet network is restricted to those devices that can authenticate using 802.1x.You can provision an AP to act as an 802.1X supplicant and authenticate to the infrastructure using the PEAP protocol.



Both Campus APs (CAPs) and Remote APs (RAPs) can be provisioned to use 802.1X authentication.

Prerequisites

- An AP has to be configured with the credentials for 802.1X authentication. These credentials are stored securely in the AP flash.
- The AP must complete the 802.1X authentication before it sends or receives IP traffic such as DHCP.



If the AP cannot complete 802.1x authentication (explicit failure or reply timeout) within 1 minute, the AP will proceed to initiate the IP traffic and attempt to contact the controller. The infrastructure can be configured to allow this. If the AP contacts the controller it will be marked as unprovisioned so that the administrator can take corrective action.

Provisioning an AP as an 802.1X Supplicant

This section describes how an AP can be provisioned as an 802.1X supplicant using CLI or the WebUI.

In the WebUI

To provision an AP as an 802.1X supplicant using the WebUI, follow these steps:

1. Navigate to the **Configuration > Wireless > AP Installation > Provisioning** window. The list of discovered APs are displayed on this page.

2. Select the AP you want to provision.
3. Click **Provision**. The provisioning window opens.
4. Select the **802.1x Parameters using PEAP** checkbox and enter the following credentials:
   a. User Name: Enter the username of the AP in the **User Name** field.
   b. Password: Enter the password of the AP in the **Password** field.
5. Enter the password again in the **Confirm Password** field and reconfirm it.
6. Click **Apply and Reboot** (at the bottom of the page).

In the CLI

To provision an AP as a 802.1X supplicant using the CLI, enter the following commands in the config mode:

```
(host) (config)# provision-ap
(host) (AP provisioning) # apdot1x-username <username>
(host) (AP provisioning) # apdot1x-passwd <password>
(host) (AP provisioning) # reprovision ap-name <apname>
```

To view the 802.1x authentication details on the controller:

```
(host) # show ap active
Active AP Table
---------------
Name       Group  IP Address   11g Clients  11g Ch/EIRP/MaxEIRP  11a Clients  11a
Ch/EIRP/MaxEIRP  AP Type  Flags  Uptime  Outer IP
----       -----  ----------   -----------  -------------------  -----------  ----------------
---  -------  -----  ------  --------
AP1X      default   10.3.15.107  0            AP:HT:1/15/21.5      0            AP:HT:44/15/21
      125      1E2    5m:48s  N/A

Flags: a = Reduce ARP packets in the air;  A = Enet1 in active/standby mode;
       B = Battery Boost On; C = Cellular; D = Disconn. Extra Calls On;
       d = Drop Mcast/Bcast On; E = Wired AP enabled; K = 802.11K Enabled;
       L = Client Balancing Enabled; M = Mesh; N = 802.11b protection disabled;
       P = PPPOE; R = Remote AP; X = Maintenance Mode;
       1 = 802.1x authenticated AP; 2 = Using IKE version 2;
```

Sample Configurations

The following examples show basic configurations on the controller for:

- Configuring Authentication with an 802.1X RADIUS Server on page 265
- Configuring Authentication with the Controller's Internal Database on page 275

In the following examples:

- Wireless clients associate to the ESSID **WLAN-01**.
- The following roles allow different networks access capabilities:
  - student
  - faculty
  - guest
  - system administrators

The examples show how to configure using the WebUI and CLI commands.

Configuring Authentication with an 802.1X RADIUS Server

- An EAP-compliant RADIUS server provides the 802.1X authentication. The RADIUS server administrator must configure the server to support this authentication. The administrator must also configure the server

to all communications with the Dell controller.

- The authentication type is WPA. From the 802.1X authentication exchange, the client and the controller derive dynamic keys to encrypt data transmitted on the wireless network.
- 802.1x authentication based on PEAP with MS-CHAPv2 provides both computer and user authentication. If a user attempts to log in without the computer being authenticated first, the user is placed into a more limited "guest" user role.

  Windows domain credentials are used for computer authentication, and the user's Windows login and password are used for user authentication. A single user sign-on facilitates both authentication to the wireless network and access to the Windows server resources.



802.1X Configuration for IAS and Windows Clients on page 1060 describes how to configure the Microsoft Internet Authentication Server and Windows XP wireless client to operate with the controller configuration shown in this section.

Configuring Roles and Policies

You can create the following policies and user roles for:

- Student
- Faculty
- Guest
- Sysadmin
- Computer

Creating the Student Role and Policy

The **student** policy prevents students from using telnet, POP3, FTP, SMTP, SNMP, or SSH to the wired portion of the network. The **student** policy is mapped to the **student** user role.

In the WebUI

1. Navigate to the **Configuration > Security > Access Control > Policies** page. Select **Add** to add the student policy.
2. For Policy Name, enter **student**.
3. For Policy Type, select **IPv4 Session**.
4. Under Rules, select **Add** to add rules for the policy.
   a. Under Source, select **user**.
   b. Under Destination, select **alias**.



The following step defines an alias representing all internal network addresses. Once defined, you can use the alias for other rules and policies.

   c. Under the alias selection, click **New**. For Destination Name, enter Internal Network. Click **Add** to add a rule. For Rule Type, select **network**. For IP Address, enter 10.0.0.0. For Network Mask/Range, enter 255.0.0.0. Click **Add** to add the network range. Repeat these steps to add the network range 172.16.0.0 - 255.255.0.0. Click **Done**. The alias Internal Network appears in the Destination menu. This step defines an alias representing all internal network addresses. Once defined, you can use the alias for other rules and policies.
   d. Under Destination, select Internal Network.
   e. Under Service, select **service**. In the Service scrolling list, select **svc-telnet**.
   f. Under Action, select **drop**.

g. Click **Add**.

5. Under Rules, click **Add**.
   a. Under Source, select **user**.
   b. Under Destination, select **alias** and then select **Internal Network**.
   c. Under Service, select **service**. In the Service scrolling list, select **svc-pop3**.
   d. Under Action, select **drop**.
   e. Click **Add**.
6. Repeat steps 4A-E to create rules for the following services: svc-ftp, svc-smtp, svc-snmp, and svc-ssh.
7. Click **Apply**.
8. Click the **User Roles** tab. Click **Add** to create the student role.
9. For Role Name, enter **student**.
10. Under Firewall Policies, click **Add**. In Choose from Configured Policies, select the student policy you previously created. Click **Done**.
11. Click **Apply**.

In the CLI

Use the following commands to configure a student role and policy:

```
(host)(config) #ip access-list session student
      user alias "Internal Network" svc-telnet deny
      user alias "Internal Network" svc-pop3 deny
      user alias "Internal Network" svc-ftp deny
      user alias "Internal Network" svc-smtp deny
      user alias "Internal Network" svc-snmp deny
      user alias "Internal Network" svc-ssh deny

(host)(config) #user-role student
      session-acl student
      session-acl allowall
```

Creating the Faculty Role and Policy

The **faculty** policy is similar to the **student** policy, however faculty members are allowed to use POP3 and SMTP for VPN remote access from home. (Students are not permitted to use VPN remote access.) The **faculty** policy is mapped to the **faculty** user role.

Using the WebUI

1. Navigate to the **Configuration > Security > Access Control > Policies** page. Click **Add** to add the faculty policy.
2. For Policy Name, enter **faculty**.
3. For Policy Type, select **IPv4 Session**.
4. Under Rules, click **Add** to add rules for the policy.
   a. Under Source, select **user**.
   b. Under Destination, select **alias**, then select **Internal Network**.
   c. Under Service, select **service**. In the Service scrolling list, select **svc-telnet**.
   d. Under Action, select **drop**.
   e. Click **Add**.
   f. Repeat steps A-E to create rules for the following services: svc-ftp, svc-snmp, and svc-ssh.
5. Click **Apply**.
6. Select the **User Roles** tab. Click **Add** to create the faculty role.

7. For Role Name, enter **faculty**.
8. Under **Firewall Policies**, click **Add**. In Choose from Configured Policies, select the faculty policy you previously created. Click **Done**.

In the CLI

Use the following commands to configure a faculty role and policy:

```
(host)(config) #ip access-list session faculty
      user alias "Internal Network" svc-telnet deny
      user alias "Internal Network" svc-ftp deny
      user alias "Internal Network" svc-snmp deny
      user alias "Internal Network" svc-ssh deny

(host)(config) #user-role faculty
      session-acl faculty
      session-acl allowall
```

Creating the Guest Role and Policy

The **guest** policy permits only access to the internet (via HTTP or HTTPS) and only during daytime working hours. The **guest** policy is mapped to the **guest** user role.

In the WebUI

1. Navigate to the **Configuration > Security > Access Control > Time Ranges** page to define the time range working-hours. Click **Add.**
   a. For Name, enter **working-hours**.
   b. For Type, select **Periodic**.
   c. Click **Add**.
   d. For Start Day, click **Weekday**.
   e. For Start Time, enter **07:30**.
   f. For End Time, enter **17:00**.
   g. Click **Done**.
   h. Click **Apply**.
2. Click the **Policies** tab. Click **Add** to add the guest policy.
3. For e**Policy Name**, enter **guest**.
4. For **Policy Type**, select **IPv4 Session**.
5. Under Rules, click **Add** to add rules for the policy.
   To create rules to permit access to DHCP and DNS servers during working hours:
   a. Under **Source**, select **user**.
   b. Under **Destination**, select **host**. In Host IP, enter **10.1.1.25**.
   c. Under **Service**, select **service**. In the Service scrolling list, select **svc-dhcp**.
   d. Under **Action**, select **permit**.
   e. Under **Time Range**, select **working-hours**.
   f. Click **Add**.
   g. Repeat steps A-F to create a rule for *svc-dns*.
   To create a rule to deny access to the internal network:
   a. Under Source, select **user**.
   b. Under Destination, select **alias**. Select **Internal Network**.
   c. Under Service, select **any**.

d. Under Action, select **drop**.
e. Click **Add**.

To create rules to permit HTTP and HTTPS access during working hours:

a. Under Source, select **user**.
b. Under Destination, select **any**.
c. Under Service, select service. In the Services scrolling list, select **svc-http**.
d. Under Action, select **permit**.
e. Under Time Range, select **working-hours**.
f. Click **Add**.
g. Repeat steps A-F for the *svc-https* service.

To create a rule that denies the user access to all destinations and all services:

a. Under Source, select **user**.
b. Under Destination, select **any**.
c. Under Service, select **any**.
d. Under Action, select **drop**.
e. Click **Add**.

6. Click **Apply**.
7. Click the **User Roles** tab. Click **Add** to create the guest role.
8. For Role Name, enter **guest.**
9. Under **Firewall Policies**, click **Add.** In Choose from Configured Policies, select the guest policy you previously created. Click **Done**.

In the CLI

Use the following commands to configure a guest role and policy:

```
   time-range working-hours periodic
      weekday 07:30 to 17:00

(host)(config) #ip access-list session guest
      user host 10.1.1.25 svc-dhcp permit time-range working-hours
      user host 10.1.1.25 svc-dns permit time-range working-hours
      user alias "Internal Network" any deny
      user any svc-http permit time-range working-hours
      user any svc-https permit time-range working-hours
      user any any deny

(host)(config) #user-role guest
      session-acl guest
```

Creating Roles and Policies for Sysadmin and Computer

The **allowall** policy, a predefined policy, allows unrestricted access to the network. The **allowall** policy is mapped to both the **sysadmin** user role and the **computer** user role.

In the WebUI

1. Navigate to **Configuration > Security > Access Control > User Roles** page. Click **Add** to create the sysadmin role.
2. For Role Name, enter **sysadmin**.
3. Under Firewall Policies, click **Add**. In Choose from Configured Policies, select the predefined **allowall** policy. Click **Done**.

4. Click **Apply**.

In the CLI

```
(host)(config) #user-role sysadmin
   session-acl allowall
```

Using the WebUI to create the computer role

1. Navigate to **Configuration > Security > Access Control > User Roles** page. Click **Add** to create the computer role.
2. For Role Name, enter **computer**.
3. Under Firewall Policies, click **Add**. In Choose from Configured Policies, select the predefined **allowall** policy. Click **Done**.
4. Click **Apply.**

Using the CLI to create the computer role

Use the following command to create a computer role:

```
(host)(config) #user-role computer
   session-acl allowall
```

Creating an Alias for the Internal Network Using the CLI

Use the following commands to create an alias for the internal network:

```
(host)(config) #netdestination "Internal Network"
   network 10.0.0.0 255.0.0.0
   network 172.16.0.0 255.255.0.0
```

Configuring the RADIUS Authentication Server

Configure the RADIUS server IAS1, with IP address 10.1.1.21 and shared key. The RADIUS server is configured to sent an attribute called Class to the controller; the value of this attribute is set to either "student," "faculty," or "sysadmin" to identify the user's group. The controller uses the literal value of this attribute to determine the role name.

On the controller, you add the configured server (IAS1) into a server group. For the server group, you configure the server rule that allows the Class attribute returned by the server to set the user role.

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. In the Servers list, select **Radius Server**. In the RADIUS Server Instance list, enter **IAS1** and click **Add**.
   a. Select IAS1 to display configuration parameters for the RADIUS server.
   b. For IP Address, enter **10.1.1.21.**
   c. For Key, enter **`|*a^t%183923!.`** (You must enter the key string twice.)
   d. Click **Apply**.
3. In the Servers list, select Server Group. In the Server Group Instance list, enter **IAS** and click **Add.**
   a. Select the server group IAS to display configuration parameters for the server group.
   b. Under Servers, click **New**.
   c. From the Server Name drop-down list, select IAS1. Click **Add Server**.
4. Under Server Rules, click **New**.
   a. For Condition, enter **Class.**
   b. For Attribute, select **value-of** from the drop-down list.
   c. For Operand, select **set role**.

d. Click **Add**.

5. Click **Apply.**

In the CLI

Use the following commands to configure the RADIUS authentication server:

```
(host)(config) #aaa authentication-server radius IAS1
   host 10.1.1.21
   key |*a^t%183923!

(host)(config) #aaa server-group IAS
   auth-server IAS1
   set role condition Class value-of
```

Configuring 802.1X Authentication

An AAA profile specifies the 802.1X authentication profile and 802.1x server group to be used for authenticating clients for a WLAN. The AAA profile also specifies the default user roles for 802.1X and MAC authentication.

In the 802.1X authentication profile, configure enforcement of machine authentication before user authentication. If a user attempts to log in before machine authentication completes, the user is placed in the limited guest role.

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > L2 Authentication** page.
2. Select **802.1X Authentication Profile**.
   a. At the bottom of the **Instance** list, enter **dot1x**, then click **Add**.
   b. Select the profile name you just added.
   c. Select **Enforce Machine Authentication**.
   d. For the Machine Authentication: Default Machine Role, select **computer**.
   e. For the Machine Authentication: Default User Role, select **guest**.
   f. Click **Apply**.
3. Select the **AAA Profiles** tab.
   a. In the AAA Profiles Summary, click **Add** to add a new profile.
   b. Enter **aaa_dot1x**, then click **Add**.
   a. Select the profile name you just added.
   b. For MAC Auth Default Role, select **computer**.
   c. For 802.1x Authentication Default Role, select **faculty**.
   d. Click **Apply**.
4. In the **Profiles** list (under the aaa_dot1x profile), select **802.1x Authentication Profile**.
   a. From the drop-down list, select the **dot1x** 802.1x authentication profile you configured previously.
   b. Click **Apply**.
5. In the **Profiles** list (under the aaa_dot1x profile), select **802.1x Authentication Server Group**.
   a. From the drop-down list, select the IAS server group you created previously.
   b. Click **Apply**.

In the CLI

Use the following commands to configure an 802.1x profile:

```
(host)(config) #aaa authentication dot1x dot1x
   machine-authentication enable
   machine-authentication machine-default-role computer
   machine-authentication user-default-role guest

(host)(config) #aaa profile aaa_dot1x
   d>ot1x-default-role faculty
   mac-default-role computer
   authentication-dot1x dot1x
   d>ot1x-server-group IAS
```

Configuring VLANs

In this example, wireless clients are assigned to either VLAN 60 or 61 while guest users are assigned to VLAN 63. VLANs 60 and 61 split users into smaller IP subnetworks, improving performance by decreasing broadcast traffic. The VLANs are internal to the Dell controller only and do not extend into other parts of the wired network. The clients' default gateway is the Dell controller, which routes traffic out to the 10.1.1.0 subnetwork.

You configure the VLANs, assign IP addresses to each VLAN, and establish the "helper address" to which client DHCP requests are forwarded.

In the WebUI

1. Navigate to the **Configuration > Network > VLANs** page. Click **Add** to add VLAN 60.
   a. For **VLAN ID**, enter **60**.
   b. Click **Apply.**
   c. Repeat steps A and B to add VLANs 61 and 63.
2. To configure IP parameters for the VLANs, navigate to the **Configuration > Network > IP > IPInterfaces** page.
   a. Click **Edit** for VLAN 60.
   b. For IP Address, enter **10.1.60.1**.
   c. For Net Mask, enter **255.255.255.0**.
   d. Under DHCP Helper Address, click **Add.** Enter **10.1.1.25** and click **Add**.
   e. Click **Apply**.
3. In the IP Interfaces page, click **Edit** for VLAN 61.
   a. For IP Address, enter **10.1.61.1**.
   b. For Net Mask, enter **255.255.255.0**.
   c. Under DHCP Helper Address, click **Add**. Enter **10.1.1.25** and click **Add.**
   d. Click **Apply**.
4. In the IP Interfaces page, click **Edit** for VLAN 63.
   a. For IP Address, enter **10.1.63.1**.
   b. For Net Mask, enter **255.255.255.0.**
   c. Under DHCP Helper Address, click **Add**. Enter **10.1.1.25** and click **Add**.
   d. Click **Apply**.
5. Select the **IP Routes** tab.
   a. For Default Gateway, enter **10.1.1.254**.
   b. Click **Apply**.

In the CLI

Use the following commands to configure VLANs:

```
(host)(config) #vlan 60
(host)(config) #interface vlan 60
   ip address 10.1.60.1 255.255.255.0
   ip helper-address 10.1.1.25

(host)(config) #vlan 61
(host)(config) #interface vlan 61
   ip address 10.1.61.1 255.255.255.0
   ip helper-address 10.1.1.25

(host)(config) #vlan 63
(host)(config) #interface vlan 63
   ip address 10.1.63.1 255.255.255.0
   ip helper-address 10.1.1.25

(host)(config) #ip default-gateway 10.1.1.254
```

Configuring the WLANs

In this example, default AP parameters for the entire network are: the default ESSID is WLAN-01 and the encryption mode is TKIP. A second ESSID called “guest” has the encryption mode set to static WEP with a configured WEP key.

In this example, the non-guest clients that associate to an AP are mapped into one of two different user VLANs. The initial AP to which the client associates determines the VLAN: clients that associate to APs in the first floor of the building are mapped to VLAN 60, and clients that associate to APs in the second floor of the building are mapped to VLAN 61. Therefore, the APs in the network are segregated into two AP groups, named first-floor and second-floor. (See Creating an AP group on page 487 for information about creating AP groups.) The guest clients are mapped into VLAN 63.

Configuring the Guest WLAN

You create and configure the virtual AP profile, guest and apply the profile to each AP group. The “guest” virtual AP profile contains the SSID profile “guest” which configures static WEP with a WEP key.

In the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. In the AP Group list, click **Edit** for first-floor.
3. Under Profiles, select **Wireless LAN** and then **Virtual AP**.
4. To create the guest virtual AP:
   a. Select **NEW** from the **Add a profile** drop-down list. Enter **guest**, and click **Add.**
   b. In the **Profile Details** entry for the guest virtual AP profile, select **NEW** from the SSID profile drop-down list. A pop-up window allows you to configure the SSID profile.
   c. For the name for the SSID profile enter **guest**.
   d. For the **Network Name** for the SSID, enter **guest**.
   e. For **Network Authentication**, select **None**.
   f. For **Encryption**, select **WEP**.
   g. Enter the WEP Key.
   h. Click **Apply** to apply the SSID profile to the Virtual AP.
   i. Under **Profile Details**, click **Apply**.
5. Click on the **guest** virtual AP name in the **Profiles** list or in **Profile Details** to display configuration parameters.
   a. Ensure that you select **Virtual AP enable**.

b. For **VLAN**, select **63**.
c. Click **Apply**.

6. Navigate to the **Configuration > Wireless > AP Configuration** page.
7. In the AP Group list, click **Edit** for the second-floor.
8. In the **Profiles** list, select **Wireless LAN** and then **Virtual AP**.
9. Select **guest** from the **Add a profile** drop-down list. Click **Add**.
10. Click **Apply**.

In the CLI

Use the following commands to configure the Guest WLAN:

```
(host)(config) #wlan ssid-profile guest
  essid guest
  wepkey1 aaaaaaaaaa
  opmode static-wep

(host)(config) #wlan virtual-ap guest
  vlan 63
  ssid-profile guest

(host)(config) #ap-group first-floor
  virtual-ap guest
(host)(config) #ap-group second-floor
  virtual-ap guest
```

Configuring the Non-Guest WLANs

You create and configure the SSID profile "WLAN-01" with the ESSID "WLAN-01" and WPA TKIP encryption. You need to create and configure two virtual AP profiles: one with VLAN 60 for the first-floor AP group and the other with VLAN 61 for the second-floor AP group. Each virtual AP profile references the SSID profile "WLAN-01" and the previously-configured AAA profile aaa_dot1x.

In the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. In the AP Group list, click **Edit** for the first-floor.
3. In the **Profiles** list, select **Wireless LAN** and then **Virtual AP**.
4. To configure the WLAN-01_first-floor virtual AP:
   a. Select **NEW** from the Add a profile drop-down list. Enter **WLAN-01_first-floor**, and click **Add**.
   b. In the **Profile Details** entry for the WLAN-01_first-floor virtual AP profile, select the **aaa_dot1x** AAA profile you previously configured. A pop-up window displays the configured AAA profile parameters. Click **Apply**.
   c. From the SSID profile drop-down list, select NEW. A pop-up window allows you to configure the SSID profile.
   d. Enter **WLAN-01** for the name of the SSID profile.
   e. For **Network Name**, enter **WLAN-01**.
   f. For **Network Authentication**, select **WPA.**
   g. Click **Apply**.
   h. At the bottom of the **Profile Details** page, click **Apply**.
5. Click on the WLAN-01_first-floor virtual AP name in the **Profiles** list or in **Profile Details** to display configuration parameters.

a. Ensurer that you select **Virtual AP enable**.
b. For **VLAN**, select 60.
c. Click **Apply**.

6. Navigate to the **Configuration > Wireless > AP Configuration** page.
7. In the **AP Group** list, click **Edit** for the second-floor.
8. In the **Profiles** list, select **Wireless LAN** and then **Virtual AP**.
9. To configure the WLAN-01_second-floor virtual AP:
   a. Select **NEW** from the **Add a profile** drop-down list. Enter **WLAN-second-floor**, and click **Add**.
   b. In the Profile Details entry for the virtual AP profile, select **aaa_dot1x** from the **AAA profile** drop-down list. A pop-up window displays the configured AAA profile parameters. Click **Apply** .
   c. From the SSID profile drop-down list, select **WLAN-01**. A pop-up window displays the configured SSID profile parameters. Click **Apply** .
   d. At the bottom of the **Profile Details** page, click **Apply.**
10. Click on the new virtual AP name in the **Profiles** list or in **Profile Details** to display configuration parameters.
    a. Ensure that you select **Virtual AP enable**.
    b. For **VLAN**, select 61.
    c. Click **Apply**.

In the CLI

Use the following commands to configure non-guest WLANs:

```
(host)(config) #wlan ssid-profile WLAN-01
   essid WLAN-01
   opmode wpa-tkip

(host)(config) #wlan virtual-ap WLAN-01_first-floor
   vlan 60
   aaa-profile aaa_dot1x
   ssid-profile WLAN-01

(host)(config) #wlan virtual-ap WLAN-01_second-floor
   vlan 61
   aaa-profile aaa_dot1x
   ssid-profile WLAN-01

(host)(config) #ap-group first-floor
   virtual-ap WLAN-01_first-floor
   ap-group second-floor
   virtual-ap WLAN-01_second-floor
```

Configuring Authentication with the Controller's Internal Database

In the following example:

- The controller's internal database provides user authentication.
- The authentication type is WPA. From the 802.1x authentication exchange, the client and the controller derive dynamic keys to encrypt data transmitted on the wireless network.

Configuring the Internal Database

Configure the internal database with the username, password, and role (student, faculty, or sysadmin) for each user. There is a default **internal** server group that includes the internal database. For the internal server

group, configure a server derivation rule that assigns the role to the authenticated client.

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. In the Servers list, select **Internal DB**.
3. Under Users, click **Add User** to add users.
4. For each user, enter a username and password.
5. Select a role for each user (if a role is not specified, the default role is guest).
6. Select the expiration time for the user account in the internal database.
7. Click **Apply**.

In the CLI



Use the privileged mode in the CLI to configure users in the controller's internal database.

Use the following command to configure the internal database:

```
(host)(config) #local-userdb add username <user> password <password>
```

Configuring a Server Rule Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **Server Group** to display the Server Group list.
3. Select the **internal** server group.
4. Under **Server Rules**, click **New** to add a server derivation rule.
   a. For **Condition**, enter Role.
   b. Select **value-of** from the drop-down list.
   c. Select **Set Role** from the drop-down list.
   d. Click **Add**.
5. Click **Apply**.

Configuring a Server Rule Using the CLI

Use the following command to configure a server rule:

```
(host)(config) #aaa server-group internal
   set role condition Role value-of
```

Configuring 802.1x Authentication

An AAA profile specifies the 802.1x authentication profile and 802.1x server group to be used for authenticating clients for a WLAN. The AAA profile also specifies the default user role for 802.1x authentication.

For this example, you enable both 802.1x authentication and termination on the controller.

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > L2 Authentication** page. In the profiles list, select **802.1x Authentication Profile**.
   a. In the **Instance** list, enter **dot1x,** then click **Add**.
   b. Select the dot1x profile you just created.
   c. Select **Termination**.

NOTE

The defaults for EAP Method and Inner EAP Method are EAP-PEAP and EAP-MSCHAPv2, respectively.

d. Click **Apply**.

2. Select the **AAA Profiles** tab.
   a. In the **AAA Profiles Summary**, click **Add** to add a new profile.
   b. Enter **aaa_dot1x**, then click **Add**.
   c. Select the aaa_dot1x profile you just created.
   d. For 802.1x Authentication Default Role, select **faculty**.
   e. Click **Apply**.
3. In the **Profiles** list (under the aaa_dot1x profile you just created), select **802.1x Authentication Profile**.
   a. Select the dot1x profile from the **802.1x Authentication Profile** drop-down list.
   b. Click **Apply.**
4. In the **Profiles** list (under the aaa_dot1x profile you just created), select **802.1x Authentication Server Group**.
   a. Select the **internal** server group.
   b. Click **Apply**.

In the CLI

Use the following commands to configure 802.1x profile:

```
(host)(config) #aaa authentication dot1x dot1x
   termination enable

(host)(config) #aaa profile aaa_dot1x
   d>ot1x-default-role student
   authentication-dot1x dot1x
   d>ot1x-server-group internal
```

Configuring VLANs

In this example, wireless clients are assigned to either VLAN 60 or 61 while guest users are assigned to VLAN 63. VLANs 60 and 61 split users into smaller IP subnetworks, improving performance by decreasing broadcast traffic. The VLANs are internal to the Dell controller only and do not extend into other parts of the wired network. The clients' default gateway is the Dell controller, which routes traffic out to the 10.1.1.0 subnetwork.

You configure the VLANs, assign IP addresses to each VLAN, and establish the "helper address" to which client DHCP requests are forwarded.

In the WebUI

1. Navigate to the **Configuration > Network > VLAN** page. Click **Add** to add VLAN 60.
   a. For **VLAN ID**, enter **60**.
   b. Click **Apply**.
   c. Repeat steps A and B to add VLANs 61 and 63.
2. To configure IP parameters for the VLANs, navigate to the **Configuration > Network > IP > IP Interfaces** page.
   a. Click **Edit** for VLAN 60.
   b. For IP Address, enter **10.1.60.1**.
   c. For Net Mask, enter **255.255.255.0**.

d. Under DHCP Helper Address, click **Add.** Enter **10.1.1.25** and click **Add**.
e. Click **Apply.**

3. In the IP Interfaces page, click **Edit** for VLAN 61.
   a. For IP Address, enter **10.1.61.1**.
   b. For Net Mask, enter **255.255.255.0**.
   c. Under DHCP Helper Address, click **Add**. Enter **10.1.1.25** and click **Add**.
   d. Click **Apply**.
4. In the IP Interfaces page, click **Edit** for VLAN 63.
   a. For IP Address, enter **10.1.63.1**.
   b. For Net Mask, enter **255.255.255.0**.
   c. Under DHCP Helper Address, click **Add**. Enter **10.1.1.25** and click **Add**.
   d. Click **Apply**.
5. Select the **IP Routes** tab.
   a. For Default Gateway, enter **10.1.1.254**.
   b. Click **Apply**.

In the CLI

Use the following commands to configure VLANs:

```
(host)(config) #vlan 60
(host)(config) #interface vlan 60
  ip address 10.1.60.1 255.255.255.0
  ip helper-address 10.1.1.25

(host)(config) #vlan 61
(host)(config) #interface vlan 61
  ip address 10.1.61.1 255.255.255.0
  ip helper-address 10.1.1.25

(host)(config) #vlan 63
(host)(config) #interface vlan 63
  ip address 10.1.63.1 255.255.255.0
  ip helper-address 10.1.1.25

(host)(config) #ip default-gateway 10.1.1.254
```

Configuring WLANs

In this example, default AP parameters for the entire network are as follows: the default ESSID is WLAN-01 and the encryption mode is TKIP. A second ESSID called guest has the encryption mode set to static WEP with a configured WEP key.

In this example, the non-guest clients that associate to an AP are mapped into one of two different user VLANs. The initial AP to which the client associates determines the VLAN: clients that associate to APs in the first floor of the building are mapped to VLAN 60, and clients that associate to APs in the second floor of the building are mapped to VLAN 61. Therefore, the APs in the network are segregated into two AP groups, named first-floor and second-floor. (See Creating an AP group on page 487 for information about creating AP groups.) The guest clients are mapped into VLAN 63.

Configuring the Guest WLAN

You create and configure the virtual AP profile, guest and apply the profile to each AP group. The guest virtual AP profile contains the SSID profile, guest which configures static WEP with a WEP key.

In the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. In the **AP Group** list, select **first-floor**.
3. In the **Profiles** list, select **Wireless LAN** and then **Virtual AP**.
4. To configure the guest virtual AP:
   a. Select **NEW** from the **Add a profile** drop-down list. Enter **guest** for the name of the virtual AP profile, and click **Add**.
   b. In the **Profile Details** entry for the guest virtual AP profile, select **NEW** from the SSID profile drop-down list. A pop-up window allows you to configure the SSID profile.
   c. Enter **guest** for the name of the SSID profile.
   d. Enter **guest** for the Network Name.
   e. For Network Authentication, select **None**.
   f. For Encryption, select **WEP**.
   g. Enter the WEP key.
   h. Click **Apply**.
   i. Under **Profile Details**, click **Apply.**
5. Click on the guest virtual AP name in the **Profiles** list or in **Profile Details** to display configuration parameters.
   a. Ensure that you select **Virtual AP enable**.
   b. For **VLAN**, select **63**.
   c. Click **Apply.**
6. Navigate to the **Configuration > Wireless > AP Configuration** page.
7. In the **AP Group** list, select **second-floor**.
8. In the **Profiles** list, select **Wireless LAN** and then **Virtual AP**.
9. Select **guest** from the **Add a profile** drop-down list. Click **Add**.
10. Click **Apply**.

In the CLI

Use the following commands to configure a Guest WLAN:

```
(host)(config) #wlan ssid-profile guest
   essid guest
   wepkey1 aaaaaaaaaa
   opmode static-wep

(host)(config) #wlan virtual-ap guest
   vlan 63
   ssid-profile guest

(host)(config) #ap-group first-floor
   virtual-ap guest
(host)(config) #ap-group second-floor
   virtual-ap guest
```

Configuring the Non-Guest WLANs

You create and configure the SSID profile “WLAN-01” with the ESSID “WLAN-01” and WPA TKIP encryption. You need to create and configure two virtual AP profiles: one with VLAN 60 for the first-floor AP group and the

other with VLAN 61 for the second-floor AP group. Each virtual AP profile references the SSID profile "WLAN-01" and the previously-configured AAA profile "aaa_dot1x".

In the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. In the AP Group list, select first-floor.
3. In the Profiles list, select Wireless LAN, then select Virtual AP.
4. To configure the WLAN-01_first-floor virtual AP:
   a. Select **NEW** from the Add a profile drop-down list. Enter **WLAN-01_first-floor**, and click **Add**.
   b. In the **Profile Details** entry for the WLAN-01_first-floor virtual AP profile, select **aaa_dot1x** from the **AAA Profile** drop-down list. A pop-up window displays the configured AAA parameters. Click **Apply**.
   c. From the SSID profile drop-down list, select **NEW**. A pop-up window allows you to configure the SSID profile.
   d. Enter **WLAN-01** for the name of the SSID profile.
   e. Enter **WLAN-01** for the Network Name.
   f. Select **WPA** for Network Authentication.
   g. Click **Apply**.
   h. At the bottom of the **Profile Details** page, click **Apply**.
5. Click on the WLAN-01_first-floor virtual AP profile name in the **Profiles** list or in **Profile Details** to display configuration parameters.
   a. Ensure that you select **Virtual AP enable**.
   b. For VLAN, select 60.
   c. Click **Apply**.
6. Navigate to the **Configuration > Wireless > AP Configuration** page.
7. In the **AP Group** list, select second-floor.
8. In the **Profiles** list, select **Wireless LAN** and then **Virtual AP**.
9. To create the WLAN-01_second-floor virtual AP:
   a. Select **NEW** from the Add a profile drop-down list. Enter **WLAN-01_second-floor**, and click **Add**.
   b. In the **Profile Details** entry for the virtual AP profile, select **aaa_dot1x** from the **AAA Profile** drop-down list. A pop-up window displays the configured AAA profile parameters. Click **Apply**.
   c. From the **SSID profile** drop-down list, select **WLAN-01**. A pop-up window displays the configured SSID profile parameters. Click **Apply**.
   d. At the bottom of the **Profile Details** page, click **Apply**.
10. Click on the WLAN-01_second-floor virtual AP profile name in the **Profiles** list or in **Profile Details** to display the configuration parameters.
    a. Ensure that you select **Virtual AP enable**.
    b. For **VLAN**, select 61.
    c. Click **Apply**.

In the CLI

Use the following commands to configure a non-guest WLAN:

```
(host)(config) #wlan ssid-profile WLAN-01
  essid WLAN-01
  opmode wpa-tkip

(host)(config) #wlan virtual-ap WLAN-01_first-floor
```

```
   vlan 60
   aaa-profile aaa_dot1x
   ssid-profile WLAN-01

(host)(config) #wlan virtual-ap WLAN-01_second-floor
   vlan 61
   aaa-profile aaa_dot1x
   sid-profile WLAN-01

(host)(config) #ap-group first-floor
   virtual-ap WLAN-01_first-floor
(host)(config) #ap-group second-floor
   virtual-ap WLAN-01_second-floor
```

Configuring Mixed Authentication Modes

Use `l2-auth-fail-through` command to perform mixed authentication which includes both MAC and 802.1x authentication. When MAC authentication fails, enable the `l2-auth-fail-through` command to perform 802.1x authentication.



By default the l2-auth-fail-through command is disabled.

Table 50: *Mixed Authentication Modes*

Authentication	1	2	3	4	5	6
MAC authentication	Success	Success	Success	Fail	Fail	Fail
802.1x authentication	Success	Fail	—	Success	Fail	—
Association	dynamic-wep	No Association	static-wep	dynamic-wep	No Association	static-wep
Role Assignment	802.1x	—	MAC	802.1x	—	logon

Table 50 describes the different authentication possibilities

In the CLI

Use the following command to configure Layer 2 authentication fail-through option:

```
(host)(config) #aaa profile test
     l2-auth-fail-through
```

Performing Advanced Configuration Options for 802.1X

This section describes advanced configuration options for 802.1X authentication.

Configuring Reauthentication with Unicast Key Rotation

When enabled, unicast and multicast keys are updated after each reauthorization. It is a best practice to configure the time intervals for reauthentication, multicast key rotation, and unicast key rotation to be at least 15 minutes. Make sure these intervals are mutually prime, and the factor of the unicast key rotation interval and the multicast key rotation interval is less than the reauthentication interval.



Unicast key rotation depends upon both the AP/controller and wireless client behavior. It is known that some wireless NICs have issues with unicast key rotation.

The following is an example of the parameters you can configure for reauthentication with unicast and multicast key rotation:

- Reauthentication: Enabled
- Reauthentication Time Interval: 6011 Seconds
- Multicast Key Rotation: Enabled
- Multicast Key Rotation Time Interval: 1867 Seconds
- Unicast Key Rotation: Enabled
- Unicast Key Rotation Time Interval: 1021 Seconds

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > L2 Authentication** page.
2. Select 802.1x Authentication Profile, then select the name of the profile you want to configure.
3. Select the **Advanced** tab. Enter the following values:
   - Reauthentication Interval: 6011
   - Multicast Key Rotation Time Interval: 1867
   - Unicast Key Rotation Time Interval: 1021
   - Multicast Key Rotation: (select)
   - Unicast Key Rotation: (select)
   - Reauthentication: (select)
4. Click **Apply**.

In the CLI

```
(host)(config) #aaa authentication dot1x profile
   reauthentication
   timer reauth-period 6011
   unicast-keyrotation
   timer ukey-rotation-period 1021
   multicast-keyrotation
   timer mkey-rotation-period 1867
```

Application Single Sign-On Using L2 Authentication

This feature allows single sign-on (SSO) for different web-based applications using Layer 2 authentication information. Single sign-on for web-based application uses Security Assertion Markup Language (SAML), which happens between the web service provider and an identity provider (IDP) that the web server trusts. A request made from the client to a web server is redirected to the IDP for authentication. If the user has already been authenticated using L2 credentials, the IDP server already knows the authentication details and returns a SAML response, redirecting the client browser to the web-based application. The user enters the web-based application without needing to enter the credentials again.

Enabling application SSO using L2 network information requires configuration on the controllerand on the IDP server. The Dell ClearPass Policy Manager (CPPM) is the only IDP supported. The controllerhas been optimized to work with CPPM to provide better functionality as an IDP.

Important Points to Remember

- CPPM is the only supported IDP.
- SSO occurs after 802.1x authentication. Therefore, SSO after captive portal authentication is not supported. Roles for captive portal and SSO are mutually exclusive and, therefore, a user in the captive portal role cannot perform SSO and vice-versa.
- SSO with VIA is not supported.
- There is a limit on the number of concurrent sessions that can be serviced at a given instant. This limit is set at the webserver level using the **web-server web-max-clients** command. The default value is 320 for W-7200 Series controller platforms and 25 for other controller platforms. The maximum number of concurrent SSO sessions that can be handled is dependent on the other web services being handled and the same time.

Enabling Application SSO

Enabling application SSO using L2 authentication information requires configuration on the controller and CPPM. This feature is enabled by completing the following steps:

- Controller:
  - Configuring an SSO-IDP Profile
  - Applying an SSO Profile to a User Role
  - Selecting an IDP Certificate
- CPPM (refer to the ClearPass Policy Manager for configuration of the following procedures):
  - Add the controller's IP address as a network device
  - Add the user to the local user DB
  - Create an enforcement profile to return the Aruba vendor-specific attribute (VSA) SSO token
  - Create an IDP attribute enforcement profile
  - Create an enforcement policy binding the Aruba VSA SSO token enforcement profile
  - Create an enforcement policy binding the IDP enforcement profile
  - Create a service, allowing the respective authentication types and authentication database, and bind the Aruba VSA SSO token enforcement policy.
  - Create a service, allowing the respective authentication types and authentication database, and bind the IDP enforcement policy.
  - Configure SSO for the CPPM.

Configuring SSO IDP-Profiles

Before SSO can be enabled, you must configure an SSO profile by completing the procedure detailed below.

In the WebUI

1. Navigate to **Configuration > Advanced Services > All Profiles > Wireless LANs > SSO**.
2. Enter the name of the SSO profile and click **Add**.
3. Click on the name of the IDP profile in the **Instance** list to edit the profile.
4. Click **New**.
5. Enter the name of the IDP URL in the **URL Name** text box.
6. Enter the IDP URL into the **URL** text box.
7. Click **Add**.
8. Repeat steps 4 through 7 for each IDP URL you are adding to the SSO profile.

9. Click **Apply** when all URLs have been added.

In the CLI

```
sso idp-profile <idp profile name>
idp <urlname> <url>
```

Applying an SSO Profile to a User Role

The newly created SSO profile must be applied to any applicable user rules that require SSO. Apply the SSO profile be completing the steps below.

In the WebUI

1. Navigate to **Configuration > Security > Access Control**.
2. Select the **User Roles** tab.
3. Select the User Role that the SSO profile will be linked to and click **Edit**.
4. Under **Misc. Configuration**, select an IDP profile from the **idp profile name** drop-down menu.
5. Click **Apply**.

In the CLI

```
user-role <role name>
sso <idp profile name>
```

Selecting an IDP Certificate

An SSL certificate is needed for SSL negotiation with browser. The certificate can be imported in PKCS12 format, so that it contains the certificate and private key, or the key pair can be generated and a certificate signing request (CSR) request sent to the enterprise CA server to generate a certificate which can then be uploaded to the controller.

For information about uploading or generating a certificate, see Managing Certificates.

After a certificate is uploaded or generated, the IDP certificate must be selected.

In the WebUI

1. Navigate to **Configuration > Management > General**.
2. Under **IDP Server Certificate**, select the IDP certificate from the **Server Certificate** drop-down menu.
3. Click **Apply**.

In the CLI

```
web-server
idp-certificate <name of the certificate>
```

Chapter 12
Stateful and WISPr Authentication

ArubaOS supports stateful 802.1X authentication, stateful NTLM authentication and authentication for Wireless Internet Service Provider roaming (WISPr). Stateful authentication differs from 802.1X authentication in that the controller does not manage the authentication process directly, but monitors the authentication messages between a user and an external authentication server, and then assigns a role to that user based upon the information in those authentication messages. WISPr authentication allows clients to roam between hotspots using different ISPs.

This chapter describes the following topics:

- Working With Stateful Authentication on page 285
- Working With WISPr Authentication on page 286
- Understanding Stateful Authentication Best Practices on page 286
- Configuring Stateful 802.1X Authentication on page 286
- Configuring Stateful NTLM Authentication on page 287
- Configuring Stateful Kerberos Authentication on page 288
- Configuring WISPr Authentication on page 289

Working With Stateful Authentication

ArubaOS supports three different types of stateful authentication:

- **Stateful 802.1X authentication**: This feature allows the controller to learn the identity and role of a user connected to a third-party AP, and is useful for authenticating users to networks with APs from multiple vendors. When an 802.1X-capable access point sends an authentication request to a RADIUS server, the controller inspects this request and the associated response to learn the authentication state of the user. It then applies an identity-based user role through the Policy Enforcement Firewall.
- **Stateful Kerberos authentication**: Use stateful Kerberos authentication to configure a controller to monitor the Kerberos authentication messages between a client and a Windows authentication server. If the client successfully authenticates via an Kerberos authentication server, the controller recognizes that the client has been authenticated and assigns that client a specified user role.
- **Stateful NTLM authentication**: NT LAN Manager (NTLM) is a suite of Microsoft authentication and session security protocols. You can use stateful NTLM authentication to configure a controller to monitor the NTLM authentication messages between a client and a Windows authentication server. If the client successfully authenticates via an NTLM authentication server, the controller recognizes that the client has been authenticated and assigns that client a specified user role.

  The default Windows authentication method changed from the older NTLM protocol to the newer Kerberos protocol, starting with Windows 2000. Therefore, stateful NTLM authentication is most useful for networks with legacy, pre-Windows 2000 clients. Note also that unlike other types of authentication, all users authenticated via stateful NTLM authentication must be assigned to the user role specified in the Stateful NTLM Authentication profile. Dell's stateful NTLM authentication does not support placing users in various roles based upon group membership or other role-derivation attributes.

Working With WISPr Authentication

WISPr authentication allows a "smart client" to authenticate on the network when they roam between Wireless Internet Service Providers, even if the wireless hotspot uses an ISP for which the client may not have an account.

If you are a hotspot operator using WISPr authentication, and a client that has an account with your ISP attempts to access the Internet at your hotspot, then your ISP's WISPr AAA server authenticates that client directly, and allows the client access on the network. If, however, the client only has an account with a *partner* ISP, then your ISP's WISPr AAA server forwards that client's credentials to the partner ISP's WISPr AAA server for authentication. Once the client has been authenticated on the partner ISP, it is authenticated on your hotspot's own ISP, as per their service agreements. After your ISP sends an authentication message to the controller, the controller assigns the default WISPr user role to that client.

ArubaOS supports the following smart clients, which enable client authentication and roaming between hotspots by embedding iPass Generic Interface Specification (GIS) *redirect*, *proxy*, *authentication*, and *logoff* messages within HTLM messages to the controller.

- iPass
- Boingo
- Trustive
- weRoam
- AT&T

Understanding Stateful Authentication Best Practices

Before you can configure a stateful authentication feature, you must define a user role you want to assign to the authenticated users, and create a server group that includes a RADIUS authentication server for stateful 802.1X authentication or a Windows server for stateful NTLM authentication. For details on performing these tasks, see the following sections of this User Guide:

- Roles and Policies on page 364
- Configuring a RADIUS Server on page 226
- Configuring a Windows Server on page 235
- Configuring Server Groups on page 238

You can use the default stateful NTLM authentication and WISPr authentication profiles to manage the settings for these features, or you can create additional profiles as desired. Note, however, that unlike most other types of authentication, stateful 802.lx authentication uses only a single Stateful 802.1X profile. This profile can be enabled or disabled, but you can not configure more than one Stateful 802.1X profile.

Configuring Stateful 802.1X Authentication

When you configure 802.1X authentication for clients on non-Dell APs, you must specify the group of RADIUS servers that performs the user authentication, and select the role to be assigned to those users who successfully complete authentication. When the user logs off or shuts down the client machine, ArubaOSnote sthe deauthentication message from the RADIUS server, and changes the user's role from the specified authenticated role back to the logon role. For details on defining a RADIUS server used for stateful 802.1X authentication, see Configuring a RADIUS Server on page 226.

In the WebUI

To configure the Stateful 802.1X Authentication profile via the WebUI:

1. Navigate to the **Configuration > Security > Authentication > L2 Authentication** window.
2. In the **Profiles** list, select **Stateful 802.1X Authentication Profile**.
3. Click the **Default Role** drop-down list, and select the role assigned to stateful 802.1X authenticated users.
4. Specify the timeout period for authentication requests, from1 to 20 seconds. The default value is 10 seconds.
5. Select the **Mode** checkbox to enable stateful 802.1X authentication.

In the CLI

Use the commands below to configure stateful 802.1X authentication via the command-line interface. The first set of commands defines the RADIUS server used for 802.1X authentication, and the second set assigns that server to a server group. The third set associates that server group with the stateful 802.1X authentication profile, then sets the authentication role and timeout period.

```
(host)(config)# aaa authentication-server radius <server-name>
   acctport <port>
   authport <port>
   clone <server>
   enable
   host <ipaddr>
   key <psk>
   nas-identifier <string>
   nas-ip <ipaddr>
   retransmit <number>
   timeout <seconds>
   use-md5
   !

(host)(config)# aaa server-group group <server-group>
   auth-server <server-name>
   !

(host)(config)# aaa authentication stateful-dot1x
   server-group <server-group>
   default-role <role>
   enable
   timeout <seconds>
```

Configuring Stateful NTLM Authentication

The Stateful NTLM Authentication profile requires that you specify a server group which includes the servers performing NTLM authentication, and the role to be assigned to users who are successfully authenticated. For details on defining a windows server used for NTLM authentication, see Configuring a Windows Server on page 235.

When the user logs off or shuts down the client machine, the user remains in the authenticated role until the user ages out, that is, until the user has sent no traffic for the amount of time specified in the User Idle Timeout setting in the **Configuration > Security > Authentication > Advanced** page.

In the WebUI

To create and configure a new instance of a stateful NTLM authentication profile via the WebUI:

1. Navigate to the **Configuration > Security > Authentication > L3 Authentication**page.
2. In the **Profiles** list, expand the **Stateful NTLM Authentication Profile**.
3. To define settings for an *existing* profile, click that profile name in the profiles list.

To create and define settings for a Stateful NTLM Authentication profile, select an existing profile, then click **Save As** in the right window pane. Enter a name for the new profile in the entry field at the top of the right window pane.

4. Click the **Default Role** drop-down list, and select the role to be assigned to all users after they complete stateful NTLM authentication.
5. Specify the timeout period for authentication requests, from 1to 20 seconds. The default value is 10 seconds.
6. Select the **Mode** checkbox to enable stateful NTLM authentication.
7. Click **Apply**.
8. In the **Profiles** list, select the **Server Group** entry below the Stateful NTLM Authentication profile.
9. Click the **Server Group** drop-down list and select the group of Windows servers you want to use for stateful NTLM authentication.
10. Click **Apply**.

In the CLI

Use the commands below to configure stateful NTLM authentication via the command-line interface. The first set of commands defines the Windows server used for NTLM authentication, the second set adds that server to a server group. The third set associates that server group with the stateful NTLM authentication profile then defines the profile settings.

```
(host)(config)# aaa authentication-server windows <windows_server_name>
   host <ipaddr>
   enable
   !

(host)(config)# aaa server-group group <server-group>
   auth-server <windows_server_name>
   !

(host)(config)# aaa authentication stateful-ntlm
   default-role <role>
   enable
   server-group <server-group>
   timeout <seconds>
```

Configuring Stateful Kerberos Authentication

The Stateful Kerberos Authentication profile requires that you specify a server group which includes the Kerberos servers and the role to be assigned to users who are successfully authenticated. For details on defining a windows server used for Kerberos authentication, see Configuring a Windows Server on page 235.

When the user logs off or shuts down the client machine, the user remains in the authenticated role until the user ages out, that is, until the user has sent no traffic for the amount of time specified in the User Idle Timeout setting in the **Configuration > Security > Authentication > Advanced** page.

In the WebUI

To create and configure a new instance of a stateful Kerberos authentication profile via the WebUI:

1. Navigate to the **Configuration > Security > Authentication > L3 Authentication** page.
2. In the **Profiles** list, expand the **Stateful Kerberos Authentication Profile**.
3. To define settings for an *existing* profile, click that profile name in the **Profiles** list.

To create and define settings for a new Stateful Kerberos Authentication profile, select an existing profile, then click **Save As** in the right window pane. Enter a name for the new profile in the entry field at the top of the right window pane.

4. Click the **Default Role** drop-down list, and select the role to be assigned to all users after they complete stateful Kerberos authentication.
5. Specify the timeout period for authentication requests, from 1-20 seconds. The default value is 10 seconds.
6. Click **Apply**.
7. In the **Profiles** list, select the **Server Group** entry below the Stateful Kerberos Authentication profile.
8. Click the **Server Group** drop-down list and select the group of Windows servers you want to use for stateful Kerberos authentication.
9. Click **Apply.**

In the CLI

Use the commands below to configure stateful Kerberos authentication via the command-line interface. The first set of commands defines the server used for Kerberos authentication, the second set adds that server to a server group, and the third set of commands associates that server group with the stateful NTLM authentication profile then defines the profile settings.

```
(host)(config)# aaa authentication-server windows <windows_server_name>
   host <ipaddr>
   enable
   !

(host)(config)# aaa server-group group <server-group>
   auth-server <windows_server_name>
   !

(host)(config)# aaa authentication stateful-kerberos
   default-role <role>
   enable
   server-group <server-group>
   timeout <seconds>
```

Configuring WISPr Authentication

A WISPr authentication profile includes parameters to define RADIUS attributes, the default role for authenticated WISPr users, maximum numbers of authenticated failures, and logon wait times. The WISPr-Location-ID sent from the controller to the WISPr RADIUS server is the concatenation of the ISO Country Code, E.164 Country Code, E.164 Area Code and SSID/Zone parameters configured in this profile.

The parameters to define WISPr RADIUS attributes are specific to the RADIUS server your ISP uses for WISPr authentication; contact your ISP to determine these values. You can find a list of ISO and ITU country and area codes at the ISO and ITU websites (iso.org) and itu.int.)

In the WebUI

This section describes how to create and configure a new instance of a WISPr authentication profile in the WebUI.

1. Navigate to the **Configuration > Security > Authentication > L3 Authentication** page.
2. In the **Profiles** list, expand the **WISPr Authentication Profile**.
3. To define settings for an *existing* profile, click that profile name in the **Profiles** list.

To create and define settings for a *new* WISPr Authentication profile, select an existing profile, then click **Save As** in the right window pane. Enter a name for the new profile in the entry field. at the top of the right window pane.

4. Define values for the parameters below.

Table 51: *WISPr Authentication Profile Parameters*

Parameter	Description
Default Role	Default role assigned to users that complete WISPr authentication.
Logon wait minimum wait	If the controller's CPU utilization has surpassed the **Login wait CPU utilization threshold value**, the **Logon wait minimum wait** parameter defines the minimum number of seconds a user has to wait to retry a login attempt. Range: 1–10 seconds. Default: 5 seconds.
Logon wait maximum wait	If the controller's CPU utilization has surpassed the **Login wait CPUutilization threshold** value, the **Logon wait maximum wait** parameter defines the maximum number of seconds a user has to wait to retry a login attempt. Range: 1–10 seconds. Default: 10 seconds.
Logon wait CPU utilization threshold	Percentage of CPU utilization at which the maximum and minimum login wait times are enforced. Range: 1–100%. Default: 60%.
WISPr Location-ID ISO Country Code	The ISO Country Code section of the WISPr Location ID.
WISPr Location-ID E.164 Country Code	The E.164 Country Code section of the WISPr Location ID.
WISPr Location-ID E.164 Area Code	The E.164 Area Code section of the WISPr Location ID.
WISPr Location-ID SSID/Zone	The SSID/Zone section of the WISPr Location ID.
WISPr Operator Name	A name identifying the hotspot operator.
WISPr Location Name	A name identifying the hotspot location. If no name is defined, the parameter uses the name of the AP to which the user has associated.

5. Click **Apply**.
6. In the **Profiles** list, select the **Server Group** entry below the WISPr Authentication profile.
7. Click the **Server Group** drop-down list and select the group of RADIUS servers you want to use for WISPr authentication.
8. Click **Apply**.



A Boingo smart client uses a NAS identifier in the format <CarrierID>_<VenueID> for location identification. To support Boingo clients, you must also configure the NAS identifier parameter in the Radius server profile for the WISPr server

In the CLI

Use the CLI commands below to configure WISPr authentication. The first set of commands defines the RADIUS server used for WISPr authentication, and the second set adds that server to a server group. The third set of commands associates that server group with the WISPR authentication profile, then defines the profile settings.

```
(host)(config)# aaa authentication-server radius <rad_server_name>
```

```
   host 172.4.77.214
   key qwERtyuIOp
   enable
   nas-identifier corp_venue1
   !

(host)(config)# aaa server-group group <server-group>
   auth-server <radius_server_name>
   !

(host)(config)# aaa authentication wispr
   default-role <role>
   logon-wait {cpu-threshold|maximum-delay|minimum-delay}
   server-group <server-group>
   wispr-location-id-ac <wispr-location-id-ac>
   wispr-location-id-cc <wispr-location-id-cc>
   wispr-location-id-isocc <wispr-location-id-isocc>
   wispr-location-id-network <wispr-location-id-network>
   wispr-location-name-location <wispr-location-name-location>
   wispr-location-name-operator-name <wispr-location-name-location>
```

The Certificate Revocation feature enables the controller to perform real-time certificate revocation checks using the Online Certificate Status Protocol (OCSP), or traditional certificate validation using the Certificate Revocation List (CRL) client.

Topics in this chapter include:

- Understanding OCSP and CRL on page 292
- Configuring the Controller as a CRL Client on page 295
- Configuring the Controller as an OCSP Responder on page 296
- Configuring the Controller as an OCSP Client on page 293
- Certificate Revocation Checking for SSH Pubkey Authentication on page 297

Understanding OCSP and CRL

OCSP (RFC 2560) is a standard protocol that consists of an OCSP client and an OCSP responder. This protocol determines revocation status of a given digital public-key certificate without downloading the entire CRL.

CRL is the traditional method of checking certificate validity. A CRL provides a list of certificate serial numbers that have been revoked or are no longer valid. CRLs let the verifier check the revocation status of the presented certificate while verifying it. CRLs are limited to 512 entries.

Both the Delegated Trust Model and the Direct Trust Model are supported to verify digitally signed OCSP responses. Unlike the Direct Trust Model, the Delegated Trust Model does not require the OCSP responder certificates to be explicitly available on the controller.

Configuring a Controller as OCSP and CRL Clients

The controller can act as an OCSP client and issues OCSP queries to remote OCSP responders located on the intranet or Internet. Since many applications in ArubaOS (such as IKE), use digital certificates, a protocol such as OCSP needs to be implemented for revocation.

An entity that relies on the content of a certificate (a relying party) needs to do the checking before accepting the certificate as being valid. One check verifies that the certificate has not been revoked. The OCSP client retrieves certificate revocation status from an OCSP responder. The responder may be the CA (Certificate Authority) that has issued the certificate in question, or it may be some other designated entity which provides the service on behalf of the CA. A *revocation checkpoint* is a logical profile that is tied to each CA certificate that the controller has (trusted or intermediate). Also, the user can specify revocation preferences within each profile.

The OCSP request is not signed by the Dell OCSP client at this time. However, the OCSP response is always signed by the responder.

Both OCSP and CRL configuration and administration is usually performed by the administrator who manages the web access policy for an organization.

In small networks where there are is no Internet connection or connection to an OCSP responder, CRL is preferable to than OCSP.

Configuring an OCSP Controller as a Responder

The controller can be configured to act as an OCSP responder (server) and respond to OCSP queries from clients that want to obtain revocation status of certificates.

The OCSP responder on the controller is accessible over HTTP port 8084. You cannot configure this port. Although the OCSP responder accepts signed OCSP requests, it does not attempt to verify the signature before processing the request. Therefore, even unsigned OCSP requests are supported.

The controller as an OCSP responder provides revocation status information to Dell applications that use CRLs. This is useful in small disconnected networks where clients cannot reach outside OCSP server to validate certificates. Typical scenarios include client to client or client to other server communication situations where the certificates of either party need to be validated.

Configuring the Controller as an OCSP Client

When OCSP is used as the revocation method, you need to configure the OCSP responder certificate and the OCSP URL.

In the WebUI

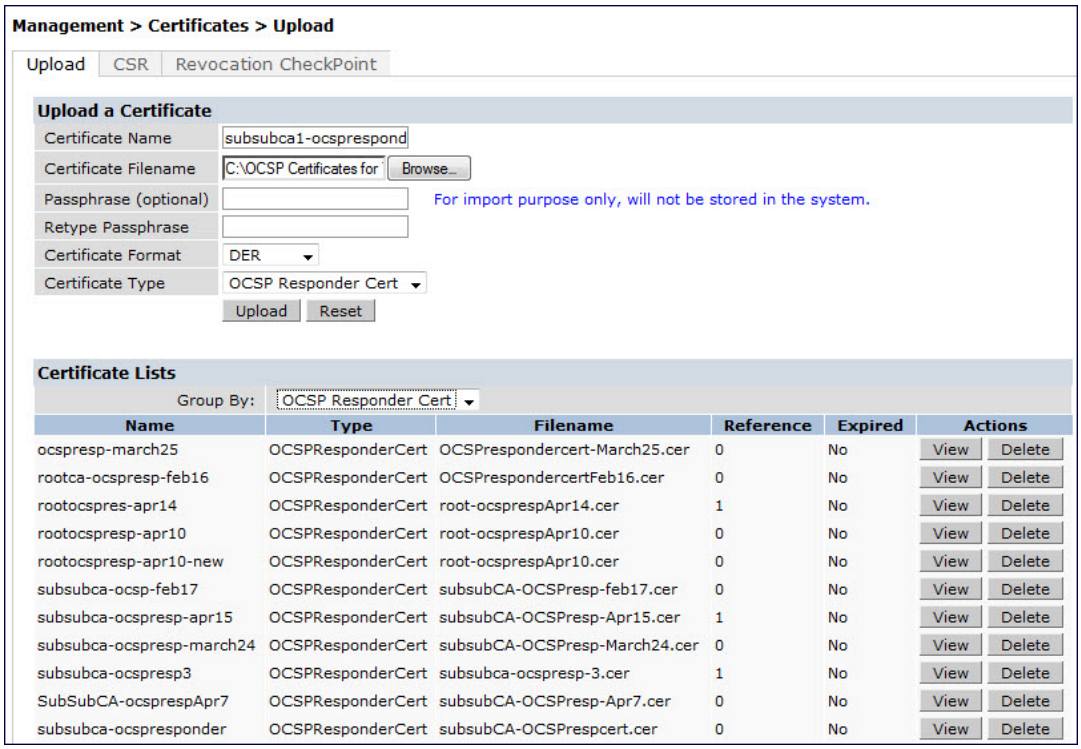
1. Navigate to the **Configuration > Management > Certificates > Upload page**.
2. Enter a name in the **Certificate Name** field. This name identifies the certificate you are uploading.
3. Enter the certificate file name in the **Certificate Filename** field. Use the **Browse** button to enter the full pathname.
4. Select the certificate format from the **Certificate Format** drop-down menu.
5. Select **OCSP Responder Cert** from the **Certificate Type** drop-down menu.



A revocation check method (OCSP or CRL) can be chosen independently for every revocation checkpoint. In this example, we are only describing the OCSP check method.

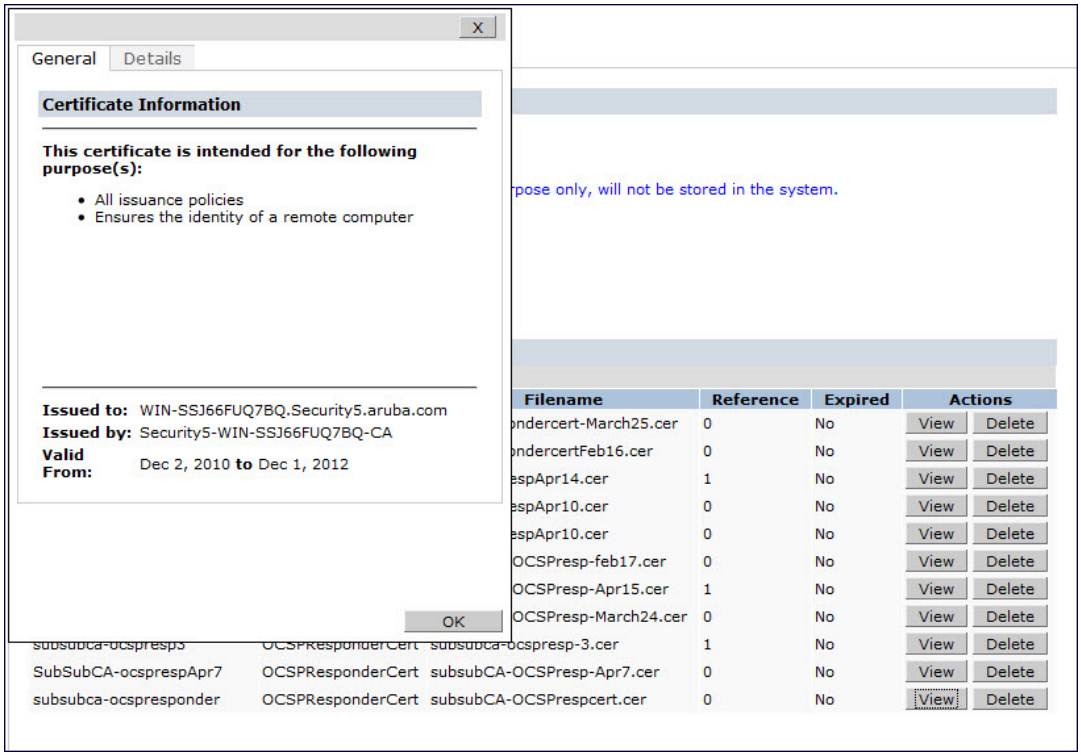
Once this certificate is uploaded it is maintained in the certificate store for OCSP responder certificates. These certificates are used for signature verification.

Figure 32 *Upload a certificate*



6. Click **Upload**. The certificate appears in the Certificate Lists pane.
7. For detailed information about an uploaded certificate, click **View** next to the certificate.

Figure 33 *View certificate details*



8. Select the **Revocation Checkpoint** tab.

9. In the **Revocation Checkpoint** pane, click Edit next to the revocation checkpoint that you want to configure. The **Revocation Checkpoint** pane displays.
10. In the **Revocation Check** field, select **ocsp** from the **Method 1** drop-down list as the primary check method.
11. In the **OCSP URL** field, enter the URL of the OCSP responder.
12. In the **OCSP Responder Cert** field, select the OCSP certificate you want to configure from the drop-down menu.
13. Click **Apply**.

In the CLI

This example configures an OCSP client with the revocation check method as OCSP for revocation check point CAroot.

The OCSP responder certificate is configured as RootCA-Ocsp_responder. The corresponding OCSP responder service is available at `http://10.4.46.202/ocsp`. The check method is OCSP for revocation check point CARoot.

```
(host) (config) #crypto-local pki rcp CARoot
   (host) (RCP-CARoot) #ocsp-responder-cert RootCA-Ocsp_responder
   (host) (RCP-CARoot) #ocsp-url http://10.4.46.202/ocsp
   (host) (RCP-CARoot) #revocation-check ocsp
```

The `show crypto-local pki OCSPResponderCert` CLI command lists the contents of the OCSP Responder Certificate store.

The `show crypto-local pki revocation checkpoint rcp_name` CLI command shows the entire configuration for a given revocation checkpoint.

Configuring the Controller as a CRL Client

CRL is the traditional method of checking certificate validity. When you want to check certificate validity using a CRL, you need to import the CRL. You can import CRLs only through the WebUI.

In the WebUI

1. Navigate to the **Configuration > Management > Certificates > Upload page**.
2. Enter a name in the **Certificate Name** field. This name identifies the CRL certificate you are uploading.
3. Enter the certificate file name in the **Certificate Filename** field. Use **Browse** to enter the full pathname.
4. Select the certificate format from the **Certificate Format** drop-down menu.
5. Select **CRL** from the **Certificate Type** drop-down menu.



A revocation check method (OCSP or CRL) can be chosen independently for every revocation checkpoint. In this example, we are only describing the CRL check method.

Once this CRL is uploaded it is maintained in the store for CRLs. These CRLs are used for signature verification.

6. Click **Upload**. The CRL appears in the Certificate Lists pane. Select **CRL** from the **Group** drop-down list if you want to display only CRLs.
7. For detailed information about an uploaded CRL, click **View** next to the CRL.
8. Select the **Revocation Checkpoint** tab.

9. In the **Revocation Checkpoint** pane, click **Edit** next to the revocation checkpoint that you want to configure. The **Revocation Checkpoint** pane displays.
10. In the **Revocation Check** field, select **crl** from the **Method 1** drop-down list.
11. In the **CRL Location** field, enter the CRL you want to use for this revocation checkpoint. The CRLs listed are files that have already been imported onto the controller.
12. Click **Apply**.

In the CLI

This example configures an OCSP responder with the check method as CRL for revocation check point ROOTCa-ssh-webui. The CRL location is crl1 and the revocation check method is crl.

```
(host) (config) #crypto-local pki rcp ROOTCa-ssh-webui
   (host) (RCP-CARoot) #crl-location file crl1
   (host) (RCP-CARoot) #revocation-check crl
```

Configuring the Controller as an OCSP Responder

When configured as an OCSP responder, the controller provides revocation status information to ArubaOS applications that use CRLs.

In the WebUI

1. Navigate to the **Configuration > Management > Certificates > Upload page**.
2. Enter a name in the **Certificate Name** field. This name identifies the OCSP signer certificate you are uploading.
3. Enter the certificate file name in the **Certificate Filename** field. Use **Browse** to enter the full pathname.
4. Select the certificate format from the **Certificate Format** drop-down menu.
5. Select **OCSP signer cert** from the **Certificate Type** drop-down menu. Once this certificate is uploaded, it is maintained in the certificate store for OCSP signer certificates. These certificates are used for signature verification.

   The OCSP signer cert signs OCSP responses for this revocation check point. The OCSP signer cert can be the same trusted CA as the check point, a designated OCSP signer certificate issued by the same CA as the check point or some other local trusted authority.

   If you do not specify an OCSP signer cert, OCSP responses are signed using the global OCSP signer certificate. If that is not present, than an error message is sent out to clients.



The OCSP signer certificate takes precedence over the global OCSP signer certificate as it is check point specific.

6. Click **Upload**. The certificate appears in the Certificate Lists pane. Select **OCSP signer cert** from the **Group** drop-down list if you want to display only those certificates which are OCSP signer certificates.
7. For detailed information about an uploaded certificate, click **View** next to the certificate.
8. Select the **Revocation Checkpoint** tab.
9. Select **Enable** next to **Enable OCSP Responder**.

   Enable OCSP Responder is a global knob that turns the OCSP responder service on or off on the controller. The default is disabled (off). Enabling this option automatically adds the OCSP responder port (TCP 8084) to the permit list in the CP firewall so this can be accessed from outside the controller.
10. Select the **OCSP signer cert** from the **OCSP Certificates** drop-down menu to be used to sign OCSP responses for this revocation check point.

11. In the **Revocation Checkpoint** pane, click **Edit** next to the revocation checkpoint that you want to configure. The **Revocation Checkpoint** pane displays.
12. In the **Revocation Check** field, optionally select a check method from the **Method 1** drop-down list. Optionally, select a backup check method from the Method 2 drop-down list.
13. Select **Enable** next to **Enable OCSP Responder**.
14. Select **OCSP signer cert** from the **OCSP Signer Cert** drop-down menu.
15. In the **CRL Location** field, enter the CRL you want used for this revocation checkpoint. The CRLs listed are files that have already been imported onto the controller.
16. Click **Apply**.

In the CLI

This example configures the controller as an OCSP responder. The OCSP responder service is enabled, the revocation check point is CAroot, the OCSP signer cert is "oscap_CA1," and the CRL file location is "Sec1-WIN-05PRGNGEKAO-CA-unrevoked.crl."

```
(host) (config) #crypto-local pki service-ocsp-responder
(host) (config) #crypto-local pki rcp CAroot
   (host) (CAroot) #ocsp-signer-cert oscsp_CA1
   (host) (CAroot) #crl-location file Sec1-WIN-05PRGNGEKAO-CA-unrevoked.crl
   (host) (CAroot) #enable-ocsp-responder
```

Certificate Revocation Checking for SSH Pubkey Authentication

This feature allows the ssh-pubkey management user to be optionally configured with a Revocation Checkpoint (RCP). This meets the requirement for a two-factor authentication and integration of device management with PKI for SSH pubkey authentication. The ArubaOS implementation of SSH using Pubkey authentication is designed for integration with smart cards or other technologies that use X.509 certificates.

The RCP checks the revocation status of the SSH user's client certificate before permitting access. If the revocation check fails, the user is denied access using the ssh-pubkey authentication method. However, the user can still authenticate through a username and password if configured to do so.

For information about configuring a revocation checkpoint, see Certificate Revocation.

Configuring the SSH Pubkey User with RCP

You can configure the SSH pubkey user with RCP to check the validity of the user's x.509 certificate.

In the WebUI

1. Navigate to **Configuration > Management > Administration**.
2. Under **Management Users**, click **Add**. The **Add User** page displays.
3. Select **Certificate Management**, then **SSH Public Key**.
4. When adding an ssh-pubkey user, when revocation check is enabled, perform either of the following tasks :
   - To enable the RCP check, select a valid configured RCP from **Revocation Checkpoint** drop-down menu.
   - Select **None** if you do not want the RCP check enabled for the ssh pubkey user.

In the CLI

The CLI allows you to configure an optional RCP for an ssh-pubkey user. Users can still be configured without the RCP. In this example, the certificate name is
"client1-rg,", the username is "test1," the role name is "root," and the rcp is "ca-rg:"

```
(host)(config) #mgmt-user ssh-pubkey client-cert client1-rg test1 root ?
rcp Revocation Checkpoint for ssh user's client certificate
```

```
(host)(config) #mgmt-user ssh-pubkey client-cert  client1-rg test1 root rcp ca-rg
```

In this example, a user is configured without the RCP:

```
(host)(config) #mgmt-user ssh-pubkey client-cert client2-rg test2 root
```

Displaying Revocation Checkpoint for the SSH Pubkey User

The RCP checks the revocation status of the SSH user's client certificate before permitting access. If the revocation check fails, the user is denied access using the ssh-pubkey authentication method. However, the user can still authenticate through a username and password if configured to do so. This feature allows the ssh-pubkey management user to be optionally configured with a Revocation Checkpoint (RCP). This meets the requirement for a two-factor authentication and integration of device management with PKI for SSH pubkey authentication. The ArubaOS implementation of SSH using Pubkey authentication is designed for integration with smart cards or other technologies that use X.50.

Configuring the SSH Pubkey User with RCP

In the WebUI

Navigate to **Configuration > Management > Administration**. The column **SSH Revocation Checkpoint** displays the RCP configured (if any) for the ssh pubkey user.

In the CLI

The **show mgmt-user ssh-pubkey** CLI command displays the column **REVOCATION CHECKPOINT** which displays the configured RCP for the ssh-pubkey user. If no RCP is configured for the user, the word none is displayed.

```
(host)#show mgmt-user ssh-pubkey
```

Removing the SSH Pubkey User

In the WebUI

1. Navigate to **Configuration > Management > Administration**.
2. Click **Delete** next to the management user you want to delete.

In the CLI

Remove ssh pubkey users by using the following command:

```
(host) (config) #no mgmt-user ssh-pubkey client-cert <certname> <username>
```

Captive portal is one of the methods of authentication supported by ArubaOS. A captive portal presents a web page which requires user action before network access is granted. The required action can be simply viewing and agreeing to an acceptable use policy, or entering a user ID and password which must be validated against a database of authorized users.

You can also configure captive portal to allow clients to download the Dell VPN dialer for Microsoft VPN clients if the VPN is to be terminated on the controller. For more information about the VPN dialer, see Virtual Private Networks on page 337.

Topics in this chapter include:

- Understanding Captive Portal on page 299
- Configuring Captive Portal in the Base Operating System on page 300
- Using Captive Portal with a PEFNG License on page 302
- Sample Authentication with Captive Portal on page 305
- Configuring Guest VLANs on page 311
- Configuring Captive Portal Authentication Profiles on page 312
- Enabling Optional Captive Portal Configurations on page 317
- Personalizing the Captive Portal Page on page 321
- Creating and Installing an Internal Captive Portal on page 323
- Creating Walled Garden Access on page 332
- Enabling Captive Portal Enhancements

Understanding Captive Portal

You can configure captive portal for guest users, where no authentication is required, or for registered users who must be authenticated against an external server or the controller's internal database.



While you can use captive portal to authenticate users, it does not provide for encryption of user data and should not be used in networks where data security is required. Captive portal is most often used for guest access, access to open systems (such as public hot spots), or as a way to connect to a VPN.

You can use captive portal for guest and registered users at the same time. The default captive portal web page provided with ArubaOS displays login prompts for both registered users and guests. (You can customize the default captive portal page, as described in Personalizing the Captive Portal Page on page 321)

You can also load up to 16 different customized login pages into the controller. The login page displayed is based on the SSID to which the client associates.

Policy Enforcement Firewall Next Generation (PEFNG) License

You can use captive portal with or without the PEFNG license installed in the controller. The PEFNG license provides identity-based security to wired and wireless clients through user roles and firewall rules. You must purchase and install the PEFNG license on the controller to use identity-based security features.

There are differences in how captive portal functions work and how you configure captive portal, depending on whether the license is installed. Other parts of this *chapter* describe how to configure captive portal in the base operating system (without the PEFNG license) and with the license installed.

Controller Server Certificate

The Dell controller is designed to provide secure services through the use of digital certificates. A server certificate installed in the controller verifies the authenticity of the controller for captive portal.

Dell controllers ship with a demonstration digital certificate. Until you install a customer-specific server certificate in the controller, this demonstration certificate is used by default for all secure HTTP connections such as captive portal. This certificate is included primarily for the purposes of feature demonstration and convenience and is not intended for long-term use in production networks. Users in a production environment are urged to obtain and install a certificate issued for their site or domain by a well-known certificate authority (CA). You can generate a Certificate Signing Request (CSR) on the controller to submit to a CA. For information on how to generate a CSR and how to import the CA-signed certificate into the controller, see Managing Certificates on page 795 in Management Access on page 778.

The controllercan accept wild card server certificates (CN begins with an asterisk). If a wildcard certificate is uploaded (for example, CN=*.domain.com), the asterisk in CN is replaced with 'captiveportal-login' in order to derive the Captive Portal logon page URL (captiveportal-login.domain.com).

Once you have imported a server certificate into the controller, you can select the certificate to be used with captive portal as described in the following sections.

To select a certificate for captive portal using the WebUI:

1. Navigate to the **Configuration > Management > General** page.
2. Under Captive Portal Certificate, select the name of the imported certificate from the drop-down list.
3. Click **Apply**.

To select a certificate for captive portal using the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #web-server
   captive-portal-cert <certificate>
```

To specify a different server certificate for captive portal with the CLI, use the **no** command to revert back to the default certificate *before* you specify the new certificate:

```
(host)(config) #web-server
   captive-portal-cert ServerCert1
   no captive-portal-cert
   captive-portal-cert ServerCert2
```

Configuring Captive Portal in the Base Operating System

The base operating system (ArubaOS without any licenses) allows full network access to all users who connect to an ESSID, both guest and registered users. In the base operating system, you cannot configure or customize user roles; this function is only available by installing the PEFNG license. Captive portal allows you to control or identify who has access to network resources.

When you create a captive portal profile in the base operating system, an implicit user role is automatically created with same name as the captive portal profile. This implicit user role allows only DNS and DHCP traffic between the client and network and directs all HTTP or HTTPS requests to the captive portal. You cannot directly modify the implicit user role or its rules. Upon authentication, captive portal clients are allowed full access to their assigned VLAN.



The WLAN Wizard within the ArubaOS WebUI allows for basic captive portal configuration for WLANs associated with the "default" ap-group: **Configuration > Wizards > WLAN Wizard**. Follow the steps in the workflow pane within the wizard and refer to the help tab for assistance.

What follows are the tasks for configuring captive portal in the base ArubaOS. The example server group and profile names appear inside quotation marks.

- Create the Server Group name. In this example, the server group name is "cp-srv".

  If you are configuring captive portal for registered users, configure the server(s) and create the server group. For more information about configuring authentication servers and server groups, see Authentication Servers on page 225.
- Create Captive Portal Authentication Profile. In this example, the profile name is "c-portal".

  Create and configure an instance of the captive portal authentication profile. Creating the captive portal profile automatically creates an implicit user role and ACL with the same name. Creating the profile "c-portal" creates an implicit user role called "c-portal". That user role allows only DNS and DHCP traffic between the client and network and directs all HTTP or HTTPS requests to the captive portal.
- Create an AAA Profile. In this example, the profile name is "aaa_c-portal".

  Create and configure an instance of the AAA profile. For the initial role, enter the implicit user role that was created in step on page 301. The initial role in the profile "aaa_c-portal" must be set to "c-portal".
- Create SSID Profile. In this example, the profile name is "ssid_c-portal".

  Create and configure an instance of the virtual AP profile which you apply to an AP group or AP name. Specify the AAA profile you created in step on page 301.
- Create a Virtual AP Profile. In this example, the profile name is "vp_c-portal".

  Create and configure an instance of the SSID profile for the virtual AP.

The following sections present the procedure for configuring the captive portal authentication profile, the AAA profile, and the virtual AP profile using the WebUI or the command line (CLI). Configuring the VLAN and authentication servers and server groups are described elsewhere in this document.

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > L3 Authentication** page. Select the **Captive Portal Authentication** profile.
   a. In the Captive Portal Authentication Profile Instance list, enter the name of the profile (for example, **c-portal**), then click **Add.**
   b. Select the captive portal authentication profile you just created.
   c. You can enable user login and/or guest login, and configure other captive portal profile parameters as described in Table 52.
   d. Click **Apply**.
2. To specify authentication servers, select Server Group under the captive portal authentication profile you just configured.
   a. Select the server group (for example, **cp-srv**) from the drop-down menu.
   b. Click **Apply.**
3. Select the **AAA Profiles** tab.
   a. In the AAA Profiles Summary, click **Add** to add a new profile. Enter the name of the profile (for example, **aaa_c-portal**), then click **Add**.
   b. Select the AAA profile you just created.

c. For Initial Role, select the captive portal authentication profile (for example, **c-portal**) you created previously.



The Initial Role must be exactly the same as the name of the captive portal authentication profile you created.

d. Click **Apply**.

4. Navigate to the **Configuration > Wireless > AP Configuration** page. Select either the AP Group or AP Specific tab. Click **Edit** for the applicable AP group name or AP name.
5. Under Profiles, select Wireless LAN, then select Virtual AP.
6. To create a new virtual AP profile, select NEW from the Add a profile drop-down menu. Enter the name for the virtual AP profile (for example, **vp_c-portal**), then click **Add**.
   a. In the Profile Details entry for the new virtual AP profile, select the AAA profile you previously created from the AAA Profile drop-down menu. A pop-up window displays the configured AAA profile parameters. Click **Apply** in the pop-up window.
   b. From the SSID profile drop-down menu, select NEW. A pop-up window allows to you configure the SSID profile.
   c. Enter the name for the SSID profile (for example, **ssid_c-portal**).
   d. Enter the Network Name for the SSID (for example, **c-portal-ap**).
   e. Click **Apply** in the pop-up window.
   f. At the bottom of the Profile Details page, click **Apply**.
7. Click on the new virtual AP name in the Profiles list or in Profile Details to display configuration parameters.
   a. Make sure Virtual AP enable is selected.
   b. For VLAN, select the VLAN to which users are assigned (for example, **20**).
   c. Click **Apply**.

In the CLI

To configure captive portal in the base operating system via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #aaa authentication captive-portal c-portal
  server-group cp-srv
(host)(config) #aaa profile aaa_c-portal
  initial-role c-portal
(host)(config) #wlan ssid-profile ssid_c-portal
  essid c-portal-ap
(host)(config) #wlan virtual-ap vp_c-portal
  aaa-profile aaa_c-portal
  ssid-profile ssid_c-portal
  vlan 20
```

Using Captive Portal with a PEFNG License

The PEFNG license provides identity-based security for wired and wireless users. There are two user roles that are important for captive portal:

- Default user role, which you specify in the captive portal authentication profile, is the role granted to clients upon captive portal authentication. This can be the predefined **guest** system role.
- Initial user role, which you specify in the AAA profile, directs clients who associate to the SSID to captive portal whenever the user initiates a Web browser connection. This can be the predefined **logon** system role.

The captive portal authentication profile specifies the captive portal login page and other configurable parameters. The initial user role configuration must include the applicable captive portal authentication profile instance.



MAC-based authentication, if enabled on the controller, takes precedence over captive portal authentication.

The following are the basic tasks for configuring captive portal using role-based access provided by the Policy Enforcement Firewall software module. Note that you must install the PEFNG license before proceeding (see Software Licenses on page 130).

- Configure the user role for a default user.

  Create and configure user roles and policies for guest or registered captive portal users. (See Roles and Policies on page 364 for more information about configuring policies and user roles.)
- Create a server group.

  If you are configuring captive portal for registered users, configure the server(s) and create the server group. (See Authentication Servers on page 225for more information about configuring authentication servers and server groups.)



If you are using the controller's internal database for user authentication, use the predefined "Internal" server group. You need to configure entries in the internal database, as described in Authentication Servers on page 225.

- Create the captive portal authentication profile.

  Create and configure an instance of the captive portal authentication profile. Specify the default user role for captive portal users.
- Configure the initial user role.

  Create and configure the initial user role for captive portal. You need to include the predefined **captiveportal** policy, which directs clients to the captive portal, in the initial user role configuration.

  You also need to specify the captive portal authentication profile instance in the initial user role configuration. For example, if you are using the predefined **logon** system role for the initial role, you need to edit the role to specify the captive portal authentication profile instance.
- Create the AAA Profile.

  Create and configure an instance of the AAA profile. Specify the initial user role.
- Create the SSID Profile "ssid_c-portal".

  Create and configure an instance of the virtual AP profile that you apply to an AP group or AP name. Specify the AAA profile you just created.
- Create the Virtual AP Profile "vp_c-portal".

  Create and configure an instance of the SSID profile for the virtual AP.

The following sections present the WebUI and Command Line (CLI) procedures for configuring the captive portal authentication profile, initial user role, the AAA profile, and the virtual AP profile. Other chapters within this document detail the configuration of the user roles and policies, authentication servers, and server groups.

Configuring Captive Portal in the WebUI

To configure captive portal with PEFNG license via the WebUI:

1. Navigate to the **Configuration > Security > Authentication > L3 Authentication** page.
2. Select Captive Portal Authentication Profile.

a. In the Captive Portal Authentication Profile Instance list, enter the name of the profile (for example, **c-portal**), then click **Add**.
b. Select the captive portal authentication profile you just created.
c. Select the default role (for example, **employee**) for captive portal users.
d. Enable guest login and/or user login, as well as other parameters (refer to Table 52).
e. Click **Apply**.

3. To specify the authentication servers, select Server Group under the captive portal authentication profile you just configured.
   a. Select the server group (for example, **cp-srv**) from the drop-down menu.
   b. Click **Apply.**
4. Select the **AAA Profiles** tab.
   a. In the AAA Profiles Summary, click **Add** to add a new profile. Enter the name of the profile (for example, **aaa_c-portal**), then click **Add**.
   b. Set the Initial role to a role that you will configure with the captive portal authentication profile.
   c. Click **Apply**.
5. Navigate to the **Configuration > Security > Access Control** page to configure the initial user role to use captive portal authentication.
   a. To edit the predefined logon role, select the **System Roles** tab, then click **Edit** for the logon role.
   b. To configure a new role, first configure policy rules in the **Policies** tab, then select the **User Roles** tab to add a new user role and assign policies.
   c. To specify the captive portal authentication profile, scroll down to the bottom of the page. Select the profile from the Captive Portal Profile drop-down menu, and click **Change**.
   d. Click **Apply**.
6. Navigate to the **Configuration > Wireless > AP Configuration** page to configure the virtual AP profile.
7. Select either the AP Group or AP Specific tab. Click **Edit** for the applicable AP group name or AP name.
8. Under Profiles, select Wireless LAN, then select Virtual AP.
9. Select NEW from the Add a profile drop-down menu to create a new virtual AP profile. Enter the name for the virtual AP profile (for example, **vp_c-portal**), then click **Add.**
   a. In the Profile Details entry for the new virtual AP profile, select the AAA profile you previously configured. A pop-up window displays the configured AAA profile parameters. Click **Apply** in the pop-up window.
   b. From the SSID profile drop-down menu, select NEW. A pop-up window allows you to configure the SSID profile.
   c. Enter the name for the SSID profile (for example, **ssid_c-portal**).
   d. Enter the Network Name for the SSID (for example, **c-portal-ap**).
   e. Click **Apply** in the pop-up window.
   f. At the bottom of the Profile Details page, click **Apply.**
10. Click on the new virtual AP name in the Profiles list or in Profile Details to display configuration parameters.
    a. Make sure Virtual AP enable is selected.
    b. For VLAN, select the VLAN to which users are assigned (for example, **20**).
    c. Click **Apply**.

Configuring Captive Portal in the CLI

To configure captive portal with the PEFNG license via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #aaa authentication captive-portal c-portal
  d>efault-role employee
  server-group cp-srv
(host)(config) #user-role logon
  captive-portal c-portal
(host)(config) #aaa profile aaa_c-portal
  initial-role logon
(host)(config) #wlan ssid-profile ssid_c-portal
  essid c-portal-ap
  vlan 20
(host)(config) #wlan virtual-ap vp_c-portal
  aaa-profile aaa_c-portal
  ssid-profile ssid_c-portal
```

Sample Authentication with Captive Portal

In the following example:

- Guest clients associate to the **guestnet** SSID which is an open wireless LAN. Guest clients are placed into VLAN 900 and assigned IP addresses by the controller's internal DHCP server. The user has no access to network resources beyond DHCP and DNS until they open a web browser and log in with a guest account using captive portal.
- Guest users are given a login and password from guest accounts created in the controller's internal database. The temporary guest accounts are created and administered by the site receptionist.
- Guest users must enter their assigned login and password into the captive portal login before they are given access to use web browsers (HTTP and HTTPS), POP3 email clients, and VPN clients (IPsec, PPTP, and L2TP) on the Internet and only during specified working hours. Guest users are prohibited from accessing internal networks and resources. All traffic to the Internet is source-NATed.



This example assumes a Policy Enforcement Firewall Next Generation (PEFNG) license is installed in the controller.

In this example, you create two user roles:

- **guest-logon** is a user role assigned to any client who associates to the guestnet SSID. Normally, any client that associates to an SSID will be placed into the *logon* system role. The **guest-logon** user role is more restrictive than the logon role.
- **auth-guest** is a user role granted to clients who successfully authenticate via the captive portal.

Creating a Guest User Role

The **guest-logon** user role consists of the following ordered policies:

- **captiveportal** is a predefined policy that allows captive portal authentication.
- **guest-logon-access** is a policy that you create with the following rules:
  - Allows DHCP exchanges between the user and the DHCP server during business hours while blocking other users from responding to DHCP requests.
  - Allows ICMP exchanges between the user and the controller during business hours.
- **block-internal-access** is a policy that you create that denies user access to the internal networks.



The **guest-logon** user role configuration needs to include the name of the captive portal authentication profile instance. You can modify the user role configuration after you create the captive portal authentication profile instance.

Creating an Auth-guest User Role

The **auth-guest** user role consists of the following ordered policies:

- **cplogout** is a predefined policy that allows captive portal logout.
- **guest-logon-access** is a policy that you create with the following rules:
  - Allows DHCP exchanges between the user and the DHCP server during business hours while blocking other users from responding to DHCP requests.
  - Allows DNS exchanges between the user and the public DNS server during business hours. Traffic is source-NATed using the IP interface of the controller for the VLAN.
- **block-internal-access** is a policy that you create that denies user access to the internal networks.
- **auth-guest-access** is a policy that you create with the following rules:
  - Allows DHCP exchanges between the user and the DHCP server during business hours while blocking other users from responding to DHCP requests.
  - Allows DNS exchanges between the user and the public DNS server during business hours. Traffic is source-NATed using the IP interface of the controller for the VLAN.
  - Allows HTTP/S traffic from the user during business hours. Traffic is source-NATed using the I interface of the controller for the VLAN.
- **drop-and-log** is a policy that you create that denies all traffic and logs the attempted network access.

Configuring Policies and Roles in the WebUI

Creating a Time Range

To create a time range via the WebUI:

1. Navigate to the **Configuration > Security > Access Control > Time Ranges** page to define the time range “working-hours”.
2. Click **Add**.
   a. For Name, enter **working-hours**.
   b. For Type, select **Periodic**.
   c. Click **Add**.
   d. For Start Day, click **Weekday**.
   e. For Start Time, enter **07:30.**
   f. For End Time, enter **17:00**.
   g. Click **Done**.
3. Click **Apply**.

To create the guest-logon-access policy via the WebUI:

1. Navigate to the **Configuration > Security > Access Control > Policies** page.
2. Select **Add** to add the guest-logon-access policy.
3. For Policy Name, enter **guest-logon-access.**
4. For Policy Type, select **IPv4 Session**.
5. Under Rules, select **Add** to add rules for the policy.
   a. Under Source, select **user**.

b. Under Destination, select **any**.
c. Under Service, select **udp**. Enter **68.**
d. Under Action, select **drop**.
e. Click **Add**.

6. Under Rules, click **Add**.
   a. Under Source, select **any**.
   b. Under Destination, select **any**.
   c. Under Service, select **service**. Select **svc-dhcp**.
   d. Under Action, select **permit.**
   e. Under Time Range, select **working-hours.**
   f. Click **Add**.

Creating Aliases

The following step defines an alias representing the public DNS server addresses. Once defined, you can use the alias for other rules and policies.

1. Navigate to the **Configuration > Security > Access Control > Policies** page.
2. Select **Add** to add the guest-logon-access policy.
3. For Policy Name, enter **guest-logon-access.**
4. For Policy Type, select **IPv4 Session**.
5. Under Rules, click **Add**.
   a. Under Source, select **user**.
   b. Under Destination, select **alias**.
   c. Under the alias selection, click **New**.
      - For Destination Name, enter "Public DNS".
      - Click **Add** to add a rule. For **Rule Type**, select **host**.
      - For **IP Address**s, enter 64.151.103.120.
      - Click **Add.** For Rule Type, select **host.**
      - For IP Address, enter 216.87.84.209.
      - Click **Add**.
      - Click **Apply.** The alias "Public DNS" appears in the Destination menu
   d. Under Destination, select Public DNS.
   e. Under Service, select **svc-dns.**
   f. Under Action, select **src-nat**.
   g. Under Time Range, select **working-hours**.
   h. Click **Add**.
6. Click **Apply**.

Creating an Auth-Guest-Access Policy

To configure the auth-guest-access policy via the WebUI:

1. Navigate to the **Configuration > Security > Access Control > Policies** page.
2. Select **Add** to add the guest-logon-access policy.
3. For Policy Name, enter **auth-guest-access**.
4. For Policy Type, select **IPv4 Session**.

5. Under Rules, select **Add** to add rules for the policy.
   a. Under Source, select **user**.
   b. Under Destination, select **any**.
   c. Under Service, select **udp**. Enter **68**.
   d. Under Action, select **drop**.
   e. Click **Add**.
6. Under Rules, click **Add**.
   a. Under Source, select **any**.
   b. Under Destination, select **any**.
   c. Under Service, select **service**. Select **svc-dhcp**.
   d. Under Action, select **permit**.
   e. Under Time Range, select **working-hours**.
   f. Click **Add**.
7. Under Rules, click **Add**.
   a. Under Source, select **user**.
   b. Under Destination, select **alias**. Select **Public DNS** from the drop-down menu.
   c. Under Service, select **service**. Select **svc-dns**.
   d. Under Action, select **src-nat**.
   e. Under Time Range, select **working-hours**.
   f. Click **Add**.
8. Under Rules, click **Add**.
   a. Under Source, select **user**.
   b. Under Destination, select **any**.
   c. Under Service, select **service**. Select **svc-http**.
   d. Under Action, select **src-nat**.
   e. Under Time Range, select **working-hours**.
   f. Click **Add**.
9. Under Rules, click **Add**.
   a. Under Source, select **user**.
   b. Under Destination, select **any**.
   c. Under Service, select **service**. Select **svc-https**.
   d. Under Action, select **src-nat**.
   e. Under Time Range, select **working-hours**.
   f. Click **Add**.
10. Click **Apply**.

Creating an Block-Internal-Access Policy

To create the block-internal-access policy via the WebUI:

1. Navigate to the **Configuration > Security > Access Control > Policies** page.
2. Select **Add** to add the block-internal-access policy.
3. For Policy Name, enter **block-internal-access.**
4. For Policy Type, select **IPv4 Session**.
5. Under Rules, select **Add** to add rules for the policy.

a. Under Source, select **user**.
b. Under Destination, select **alias**.



The following step defines an alias representing all internal network addresses. Once defined, you can use the alias for other rules and policies.

c. Under the alias selection, click **New**. For Destination Name, enter "Internal Network". Click **Add** to add a rule. For Rule Type, select **network**. For IP Address, enter 10.0.0.0. For Network Mask/Range, enter 255.0.0.0. Click **Add** to add the network range. Repeat these steps to add the network ranges 172.16.0.0 255.255.0.0 and 192.168.0.0 255.255.0.0. Click **Apply**. The alias "Internal Network" appears in the Destination menu
d. Under Destination, select Internal Network.
e. Under Service, select **any**.
f. Under Action, select **drop**.
g. Click **Add**.

6. Click **Apply**.

Creating a Drop-and-Log Policy

To create the drop-and-log policy via the WebUI:

1. Navigate to the **Configuration > Security > Access Control > Policies** page.
2. Select **Add** to add the drop-and-log policy.
3. For Policy Name, enter **drop-and-log**.
4. For Policy Type, select **IPv4 Session**.
5. Under Rules, select **Add** to add rules for the policy.
   a. Under Source, select **user.**
   b. Under Destination, select **any**.
   c. Under Service, select **any.**
   d. Under Action, select **drop**.
   e. Select **Log**.
   f. Click **Add**.
6. Click Apply.

Creating a Guest Role

To create a guest role via the WebUI:

1. Navigate to the **Configuration > Security > Access Control > User Roles** page.
2. Click **Add**.
3. For Role Name, enter guest-logon.
4. Under Firewall Policies, click **Add**.
5. For Choose from Configured Policies, select captiveportal from the drop-down menu.
6. Click **Done.**
7. Under Firewall Policies, click **Add.**
8. For Choose from Configured Policies, select guest-logon-access from the drop-down menu.
9. Click **Done**.
10. Under Firewall Policies, click **Add**.

11. For Choose from Configured Policies, select block-internal-access from the drop-down menu.
12. Click **Done**.
13. Click **Apply**.

Creating an Auth-Guest Role

To create the guest-logon role via the WebUI:

1. Navigate to the **Configuration > Security > Access Control > User Roles** page.
2. Click **Add.**
3. For Role Name, enter auth-guest.
4. Under Firewall Policies, click **Add**.
5. For Choose from Configured Policies, select cplogout from the drop-down menu.
6. Click **Done**.
7. Under Firewall Policies, click **Add**.
8. For Choose from Configured Policies, select guest-logon-access from the drop-down menu.
9. Click **Done**.
10. Under Firewall Policies, click **Add.**
11. For Choose from Configured Policies, select block-internal-access from the drop-down menu.
12. Click **Done**.
13. Under Firewall Policies, click **Add.**
14. For Choose from Configured Policies, select auth-guest-access from the drop-down menu.
15. Click **Done**.
16. Under Firewall Policies, click **Add**.
17. For Choose from Configured Policies, select drop-and-log from the drop-down menu.
18. Click **Done**.
19. Click **Apply**.

Configuring Policies and Roles in the CLI

Defining a Time Range

To create a time range via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #time-range working-hours periodic
   weekday 07:30 to 17:00
```

Creating Aliases

To create aliases via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #netdestination "Internal Network"
network 10.0.0.0 255.0.0.0
network 172.16.0.0 255.255.0.0
   network 192.168.0.0 255.255.0.0
(host)(config) #netdestination "Public DNS"
   host 64.151.103.120
   host 216.87.84.209
```

Creating a Guest-Logon-Access Policy

To create a guest-logon-access policy via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #ip access-list session guest-logon-access
   user any udp 68 deny
   any any svc-dhcp permit time-range working-hours
   user alias "Public DNS" svc-dns src-nat time-range working-hours
```

Creating an Auth-Guest-Access Policy

To create an auth-guest-access policy via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #ip access-list session auth-guest-access
   user any udp 68 deny
   any any svc-dhcp permit time-range working-hours
   user alias "Public DNS" svc-dns src-nat time-range working-hours
   user any svc-http src-nat time-range working-hours
   user any svc-https src-nat time-range working-hours
```

Creating a Block-Internal-Access Policy

To create a block-internal-access policy via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #ip access-list session block-internal-access
   user alias "Internal Network" any deny
```

Creating a Drop-and-Log Policy

To create a drop-and-log policy via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #ip access-list session drop-and-log
   user any any deny log
```

Creating a Guest-Logon Role

To create a guest-logon-role via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #user-role guest-logon
   session-acl captiveportal position 1
   session-acl guest-logon-access position 2
   session-acl block-internal-access position 3
```

Creating an Auth-Guest Role

To create an auth-guest role via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #user-role auth-guest
   session-acl cplogout position 1
   session-acl guest-logon-access position 2
   session-acl block-internal-access position 3
   session-acl auth-guest-access position 4
   session-acl drop-and-log position 5
```

Configuring Guest VLANs

Guests using the WLAN are assigned to VLAN 900 and are given IP addresses via DHCP from the controller.

In the WebUI

1. Navigate to the **Configuration > Network > VLANs** page.
   a. Select the **VLAN ID** tab.
   a. Click **Add**.
   b. For VLAN ID, enter 900.
   c. Click **Apply**.
2. Navigate to the **Configuration > Network > IP > IP Interfaces** page.
   a. Click the **IP Interfaces** tab.
   a. Click **Edit** for VLAN 900.
   b. For IP Address, enter 192.168.200.20.
   c. For Net Mask, enter 255.255.255.0.
   d. Click **Apply**.
3. Click the **DHCP Server** tab.
   a. Select **Enable DHCP Server**.
   b. Click **Add** under Pool Configuration.
   c. In the **Pool Name** field, enter **guestpool**.
   d. In the **Default Router** field, enter 192.168.200.20.
   e. In the **DNS Server** field, enter 64.151.103.120.
   f. In the **Lease** field, **enter** 4 hours.
   g. In the **Network** field, enter 192.168.200.0. In the **Netmask** field, enter 255.255.255.0.
   h. Click **Done.**
4. Click **Apply**.

In the CLI

```
(host)(config) #vlan 900
(host)(config) #interface vlan 900
(host)(config) #ip address 192.168.200.20 255.255.255.0
(host)(config) #ip dhcp pool "guestpool"
(host)(config) #default-router 192.168.200.20
(host)(config) #dns-server 64.151.103.120
(host)(config) #lease 0 4 0
(host)(config) #network 192.168.200.0 255.255.255.0
```

Configuring Captive Portal Authentication Profiles

In this section, you create an instance of the captive portal authentication profile and the AAA profile. For the captive portal authentication profile, you specify the previously-created **auth-guest** user role as the default user role for authenticated captive portal clients and the authentication server group (“Internal”).

To configure captive portal authentication via the WebUI:

1. Navigate to the **Configuration > Security > Authentication > L3 Authentication** page. In the Profiles list, select Captive Portal Authentication Profile.
   a. In the Captive Portal Authentication Profile Instance list, enter **guestnet** for the name of the profile, then click **Add**.
   b. Select the captive portal authentication profile you just created.
   c. For Default Role, select **auth-guest**.
   d. Select User Login.

e. Deselect (uncheck) Guest Login.
f. Click **Apply**.

2. Select **Server Group** under the **guestnet** captive portal authentication profile you just created.
   a. Select **internal** from the Server Group drop-down menu.
   b. Click **Apply**.

To configure captive portal authentication via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #aaa authentication captive-portal guestnet
  d>efault-role auth-guest
  user-logon
  no guest-logon
  server-group internal
```

Modifying the Initial User Role

The captive portal authentication profile specifies the captive portal login page and other configurable parameters. The initial user role configuration must include the applicable captive portal authentication profile instance. Therefore, you need to modify the **guest-logon** user role configuration to include the **guestnet** captive portal authentication profile.

To modify the guest-logon role via the WebUI:

1. Navigate to the **Configuration > Security > Access Control > User Roles** page.
2. Select **Edit** for the guest-logon role.
3. Scroll down to the bottom of the page.
4. Select the captive portal authentication profile you just created from the Captive Portal Profile drop-down menu, and click **Change.**
5. Click **Apply**.

To modify the guest-logon role via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #user-role guest-logon
  captive-portal guestnet
```

Configuring the AAA Profile

In this section, you configure the **guestnet** AAA profile, which specifies the previously-created **guest-logon** role as the initial role for clients who associate to the WLAN.

To configure the AAA profile via the WebUI:

1. Navigate to the **Configuration > Security > Authentication > AAA Profiles** page.
2. In the AAA Profiles Summary, click **Add** to add a new profile. Enter **guestnet** for the name of the profile, then click **Add**.
3. For Initial role, select guest-logon.
4. Click **Apply.**

To configure the AAA profile via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #aaa profile guestnet
  initial-role guest-logon
```

Configuring the WLAN

In this section, you create the **guestnet** virtual AP profile for the WLAN. The **guestnet** virtual AP profile contains the SSID profile **guestnet** (which configures opensystem for the SSID) and the AAA profile **guestnet**.

To configure the guest WLAN via the WebUI:

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. Select either AP Group or AP Specific tab. Click **Edit** for the AP group or AP name.
3. To configure the virtual AP profile, navigate to the **Configuration > Wireless > AP Configuration** page. Select either the AP Group or AP Specific tab. Click **Edit** for the applicable AP group name or AP name.
4. Under Profiles, select Wireless LAN, then select Virtual AP.
5. To create a new virtual AP profile, select NEW from the Add a profile drop-down menu. Enter the name for the virtual AP profile (for example, **guestnet**), and click **Add**.
   a. In the Profile Details entry for the new virtual AP profile, select the AAA profile you previously configured. A pop-up window displays the configured AAA profile parameters. Click **Apply** in the pop-up window.
   b. From the SSID profile drop-down menu, select NEW. A pop-up window allows you to configure the SSID profile.
   c. Enter the name for the SSID profile (for example, **guestnet**).
   d. Enter the Network Name for the SSID (for example, **guestnet**).
   e. For Network Authentication, select None.
   f. For Encryption, select Open.
   g. Click **Apply** in the pop-up window.
   h. At the bottom of the Profile Details page, click **Apply.**
6. Click on the new virtual AP name in the Profiles list or in Profile Details to display configuration parameters.
   a. Make sure Virtual AP enable is selected.
   b. For VLAN, select the ID of the VLAN in which captive portal users are placed (for example, VLAN **900)**.
   c. Click **Apply.**

To configure the guest WLAN via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #wlan ssid-profile guestnet
  essid guestnet
  opmode opensystem

(host)(config) #aaa profile guestnet
  initial-role guest-logon

(host)(config) #wlan virtual-ap guestnet
  vlan 900
  aaa-profile guestnet
  ssid-profile guestnet
```

Managing User Accounts

Temporary user accounts are created in the internal database on the controller. You can create a user role which will allow a receptionist to create temporary user accounts. Guests can use the accounts to log into a captive portal login page to gain Internet access.

See Creating Guest Accounts on page 813 for more information about configuring guest provisioning users and administering guest accounts.

Configuring Captive Portal Configuration Parameters

Table 52 describes configuration parameters on the WebUI Captive Portal Authentication profile page.



In the CLI, you configure these options with the **aaa authentication captive-portal** commands.

Table 52: *Captive Portal Authentication Profile Parameters*

Parameter	Description
Default Role	Role assigned to the Captive Portal user upon login. When both user and guest logon are enabled, the default role applies to the user logon; users logging in using the guest interface are assigned the guest role. Default: guest
Default Guest Role	Role assigned to guest. Default: guest
Redirect Pause	Time, in seconds, that the system remains in the initial welcome page before redirecting the user to the final web URL. If set to 0, the welcome page displays until the user clicks on the indicated link. Default: 10 seconds
Login Page	URL of the page that appears for the user logon. This can be set to any URL. Default: /auth/index.html
User Logon	Enables Captive Portal with authentication of user credentials. Default: Enabled
Guest Login	Enables Captive Portal logon without authentication. Default: Disabled
Logout popout window	Enables a pop-up window with the Logout link for the user to logout after logon. If this is disabled, the user remains logged in until the user timeout period has elapsed or the station reloads. Default: Enabled
Use HTTP for authentication	Use HTTP protocol on redirection to the Captive Portal page. If you use this option, modify the captive portal policy to allow HTTP traffic. Default: disabled (HTTPS is used)
Logon wait minimum wait	Minimum time, in seconds, the user will have to wait for the logon page to pop up if the CPU load is high. This works in conjunction with the Logon wait CPU utilization threshold parameter. Default: 5 seconds
Logon wait maximum wait	Configure parameters for the logon wait interval Default: 10 seconds
Logon wait CPU utilization threshold	CPU utilization percentage above which the Logon wait interval is applied when presenting the user with the logon page. Default: 60%
Max Authentication failures	Maximum number of authentication failures before the user is blacklisted. Default: 0
Show FDQN	Allows the user to see and select the fully-qualified domain name (FQDN) on the login page. The FQDNs shown are specified when configuring individual servers for the server group used with captive portal authentication.

Parameter	Description
	Default: Disabled
Authentication Protocol	Select the PAP, CHAP or MS-CHAPv2 authentication protocol. **NOTE:** Do not use the CHAP = option unless instructed to do so by aDell representative.
Logon Page	URL of the page that appears before logon. This can be set to any URL. Default: /auth/index.html
Welcome Page	URL of the page that appears after logon and before redirection to the web URL. This can be set to any URL. Default: /auth/welcome.html
Show Welcome Page	Displays the configured welcome page before the user is redirected to their original URL. If this option is disabled, users are redirected to the web URL immediately after they log in. Default: Enabled
Add switch IP address in redirection URL	Sends the controller's IP address in the redirection URL when external captive portal servers are used. An external captive portal server can determine the controller from which a request originated by parsing the 'switchip' variable in the URL. Default: Disabled
Add User VLAN in the Redirection URL	Sends the user's VLAN ID in the redirection URL when external captive portal servers are used.
Add a controller interface in the redirection URL	Sends the controller's interface IP address in the redirection URL when external captive portal servers are used. An external captive portal server can determine the controller from which a request originated by parsing the 'switchip' variable in the URL. This parameter requires the Public Access license.
Allow only one active user session	Allows only one active user session at a time. Default: Disabled
White List	To add a netdestination to the captive portal whitelist, enter the destination host or subnet, then click **Add.** The netdestination will be added to the whitelist. To remove a netdestination from the whitelist, select it in the whitelist field, then click **Delete**. If you have not yet defined a netdestination, use the CLI command **netdestination** to define a destination host or subnet before you add it to the whitelist. This parameter requires the Public Access license.
Black List	To add a netdestination to the captive portal blacklist, enter the destination host or subnet, then click **Add**. The netdestination will be added to the blacklist. To remove a netdestination from the blacklist, select it in the blacklist field, then click **Delete**. If you have not yet defined a netdestination, use the CLI command **netdestination** to define a destination host or subnet before you add it to the blacklist.
Show Acceptable Use Policy Page	Show the acceptable use policy page before the logon page. Default: Disabled

Parameter	Description
User idle timeout	The user idle timeout value for this profile. Specify the idle timeout value for the client in seconds. Valid range is 30-15300 in multiples of 30 seconds. Enabling this option overrides the global settings configured in the AAA timers. If this is disabled, the global settings are used.
Redirect URL	URL to which an authenticated user will be directed. This parameter must be an absolute URL that begins with either **http://** or **https://**.
URL Hash Key	If a redirection URL is defined, enter a URL Hash Key to hash the redirect URL using the specified key. This parameter enhances security for the Clearpass Guest login URL so that Clearpass can trust and ensure that the client MAC address in the redirect URL has not been tampered with by anyone. Default: Disabled.

Enabling Optional Captive Portal Configurations

The following are optional captive portal configurations:

- Uploading Captive Portal Pages by SSID Association on page 317
- Changing the Protocol to HTTP on page 318
- Configuring Redirection to a Proxy Server on page 319
- Redirecting Clients on Different VLANs on page 320
- Web Client Configuration with Proxy Script on page 320

Uploading Captive Portal Pages by SSID Association

You can upload custom login pages for captive portal into the controller through the WebUI (refer to Creating and Installing an Internal Captive Portal on page 323). The SSID to which the client associates determines the captive portal login page displayed.

You specify the captive portal login page in the captive portal authentication profile, along with other configurable parameters. The initial user role configuration must include the applicable captive portal authentication profile instance. (In the case of captive portal in the base operating system, the initial user role is automatically created when you create the captive portal authentication profile instance.) You then specify the initial user role for captive portal in the AAA profile for the WLAN.

When you have multiple captive portal login pages loaded in the controller, you must configure a unique initial user role and user role, and captive portal authentication profile, AAA profile, SSID profile, and virtual AP profile for each WLAN that will use captive portal. For example, if you want to have different captive portal login pages for the engineering, business and faculty departments, you need to create and configure according to Table 53.

Table 53: *Captive Portal login Pages*

Entity	Engineering	Business	Faculty
Captive portal login page	/auth/eng-login.html	/auth/bus-login.html	/auth/fac-login.html
Captive portal user role	eng-user	bus-user	fac-user
Captive portal authentication profile	eng-cp (Specify /auth/eng-login.html and eng-user)	bus-cp (Specify /auth/bus-login.html and bus-user)	fac-cp (Specify /auth/bus-login.html and fac-user)
Initial user role	eng-logon (Specify the eng-cp profile)	bus-logon (Specify the bus-cp profile)	fac-logon (Specify the fac-logon profile)
AAA profile	eng-aaa (Specify the eng-logon user role)	bus-aaa (Specify the bus-logon user role)	fac-aaa (Specify the fac-logon user role)
SSID profile	eng-ssid	bus-ssid	fac-ssid
Virtual AP profile	eng-vap	bus-vap	fac-vap

Changing the Protocol to HTTP

By default, the HTTPS protocol is used on redirection to the Captive Portal page. If you need to use HTTP instead, you need to do the following:

- Modify the captive portal authentication profile to enable the HTTP protocol.
- *For captive portal with role-based access only*—Modify the **captiveportal** policy to permit HTTP traffic instead of HTTPS traffic.

In the base operating system, the implicit ACL captive-portal-profile is automatically modified.

To change the protocol to HTTP via the WebUI:

1. Edit the captive portal authentication profile by navigating to the **Configuration > Security > Authentication > L3 Authentication** page.
   a. Enable (select) "Use HTTP for authentication".
   b. Click **Apply.**
2. (For captive portal with role-based access only) Edit the **captiveportal** policy by navigating to the **Configuration > Security > Access Control > Policies** page.
   a. Delete the rule for "user mswitch svc-https dst-nat".
   b. Add a new rule with the following values and move this rule to the top of the rules list:
      - source is user
      - destination is the mswitch alias
      - service is svc-http
      - action is dst-nat
   c. Click **Apply.**

To change the protocol to HTTP via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #aaa authentication captive-portal profile
  protocol-http
```

```
(For captive portal with role-based access only)
(host)(config) #ip access-list session captiveportal
   no user alias mswitch svc-https dst-nat
   user alias mswitch svc-http dst-nat
   user any svc-http dst-nat 8080
   user any svc-https dst-nat 8081
```

Configuring Redirection to a Proxy Server

You can configure captive portal to work with proxy Web servers. When proxy Web servers are used, browser proxy server settings for end users are configured for the proxy server's IP address and TCP port. When the user opens a Web browser, the HTTP/S connection request must be redirected from the proxy server to the captive portal on the controller.

To configure captive portal to work with a proxy server:

- (For captive portal with base operating system) Modify the captive portal authentication profile to specify the proxy server's IP address and TCP port.
- (For captive portal with role-based access) Modify the **captiveportal** policy to have traffic for the proxy server's port destination NATed to port 8088 on the controller.

The base operating system automatically modifies the implicit ACL *captive-portal-profile*.

The following sections describe how use the WebUI and CLI to configure the captive portal with a proxy server.



When HTTPS traffic is redirected from a proxy server to the controller, the user's browser will display a warning that the subject name on the certificate does not match the hostname to which the user is connecting.

To redirect proxy server traffic using the WebUI:

1. For captive portal with Dell base operating system, edit the captive portal authentication profile by navigating to the **Configuration > Security > Authentication > L3 Authentication** page.
   a. For Proxy Server, enter the IP address and port for the proxy server.
   b. Click **Apply**.
2. For captive portal with role-based access, edit the **captiveportal** policy by navigating to the **Configuration > Security > Access Control > Policies** page.
3. Add a new rule with the following values:
   a. Source is user
   b. Destination is any
   c. Service is TCP
   d. Port is the TCP port on the proxy server
   e. Action is dst-nat
   f. IP address is the IP address of the proxy port
   g. Port is the port on the proxy server
4. Click **Add** to add the rule. Use the up arrows to move this rule just below the rule that allows HTTP(S) traffic.
5. Click **Apply**.

To redirect proxy server traffic via the command-line interface, access the CLI in config mode and issue the following commands.

For captive portal with Dell base operating system:

```
(host)(config) #aaa authentication captive-portal profile
   proxy host ipaddr port port
```

For captive portal with role-based access:

```
(host)(config) #ip access-list session captiveportal
   user alias mswitch svc-https permit
   user any tcp port dst-nat 8088
   user any svc-http dst-nat 8080
   user any svc-https dst-nat 8081
```

Redirecting Clients on Different VLANs

You can redirect wireless clients that are on different VLANs (from the controller's IP address) to the captive portal on the controller. To do this:

1. Specify the redirect address for the captive portal.
2. For captive portal with the PEFNG license only, you need to modify the captiveportal policy that is assigned to the user. To do this:
   a. Create a network destination alias to the controller interface.
   b. Modify the rule set to allow HTTPS to the new alias instead of the mswitch alias.



In the base operating system, the implicit ACL *captive-portal-profile* is automatically modified.

This example shows how to use the command-line interface to create a network destination called cp-redirect and use that in the captiveportal policy:

```
(host)(config) #ip cp-redirect-address ipaddr
```

For captive portal with PEFNG license:

```
(host)(config) #netdestination cp-redirect ipaddr
(host)(config) #ip access-list session captiveportal
   user alias cp-redirect svc-https permit
   user any svc-http dst-nat 8080
   user any svc-https dst-nat 8081
```

Web Client Configuration with Proxy Script

If the web client proxy configuration is distributed through a proxy script (a `.pac` file), you need to configure the **captiveportal** policy to allow the client to download the file. Note that in order modify the captiveportal policy, you must have the PEFNG license installed in the controller.

To allow clients to download proxy script via the WebUI:

1. Edit the **captiveportal** policy by navigating to the **Configuration > Security > Access Control > Policies** page.
2. Add a new rule with the following values:
   - Source is user
   - Destination is host
   - Host IP is the IP address of the proxy server
   - Service is svc-https or svc-http
   - Action is permit
3. Click **Add** to add the rule. Use the up arrows to move this rule above the rules that perform destination NAT.
4. Click **Apply**.

To allow clients to download proxy script via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #ip access-list session captiveportal
   user alias mswitch svc-https permit
   user any tcp port dst-nat 8088
   user host ipaddr svc-https permit
   user any svc-http dst-nat 8080
   user any svc-https dst-nat 8081
```

Personalizing the Captive Portal Page

The following can be personalized on the default captive portal page:

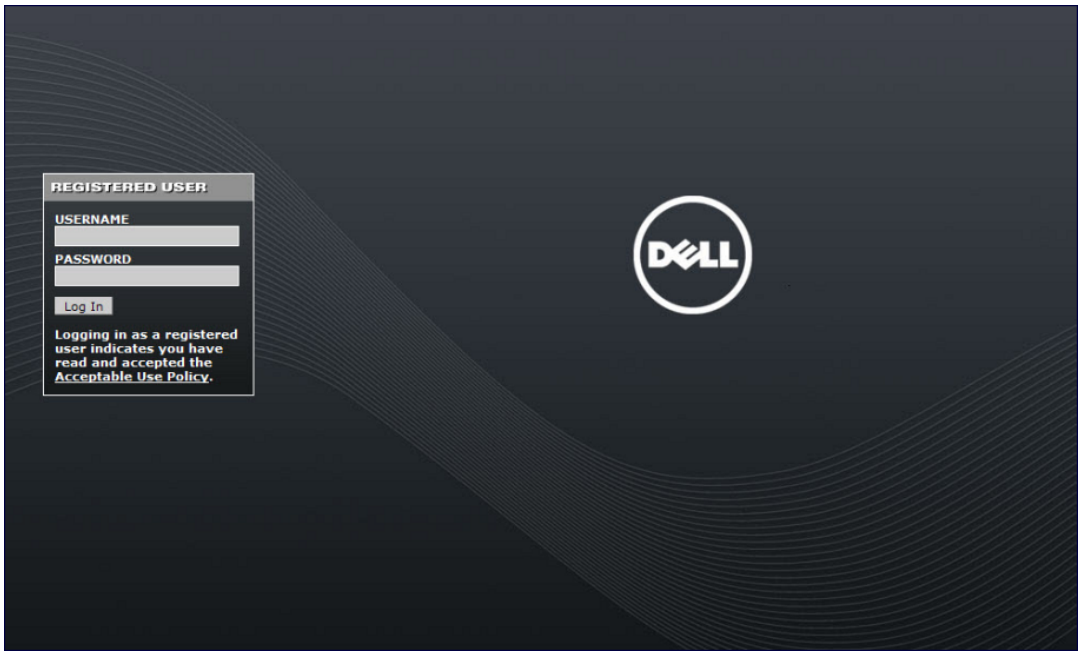
- Captive portal background
- Page text
- Acceptance Use Policy

The background image and text should be visible to users with a browser window on a 1024 by 768 pixel screen. The background should not clash if viewed on a much larger monitor. A good option is to have the background image at 800 by 600 pixels, and set the background color to be compatible. The maximum image size for the background can be around 960 by 720 pixels, as long as the image can be cropped at the bottom and right edges. Leave space on the left side for the login box.

You can create your own web pages and install them in the controller for use with captive portal. See “Internal Captive Portal” on page 265

1. Navigate to the **Configuration > Management > Captive Portal > Customize Login Page** page.

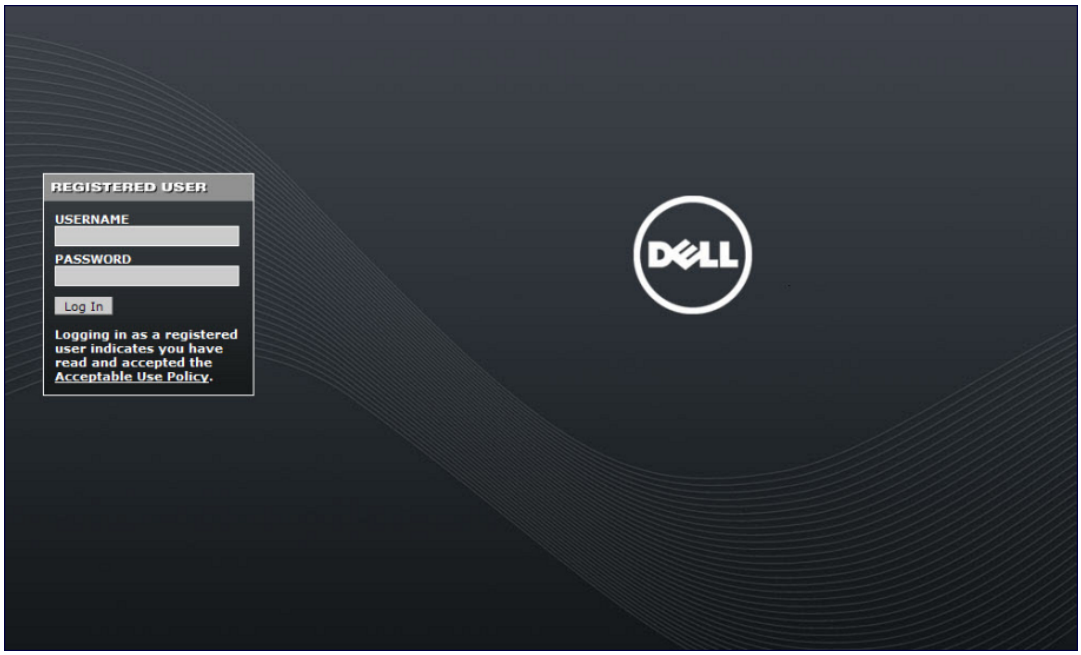
   You can choose one of three page designs. To select an existing design, click the first or the second page design present.



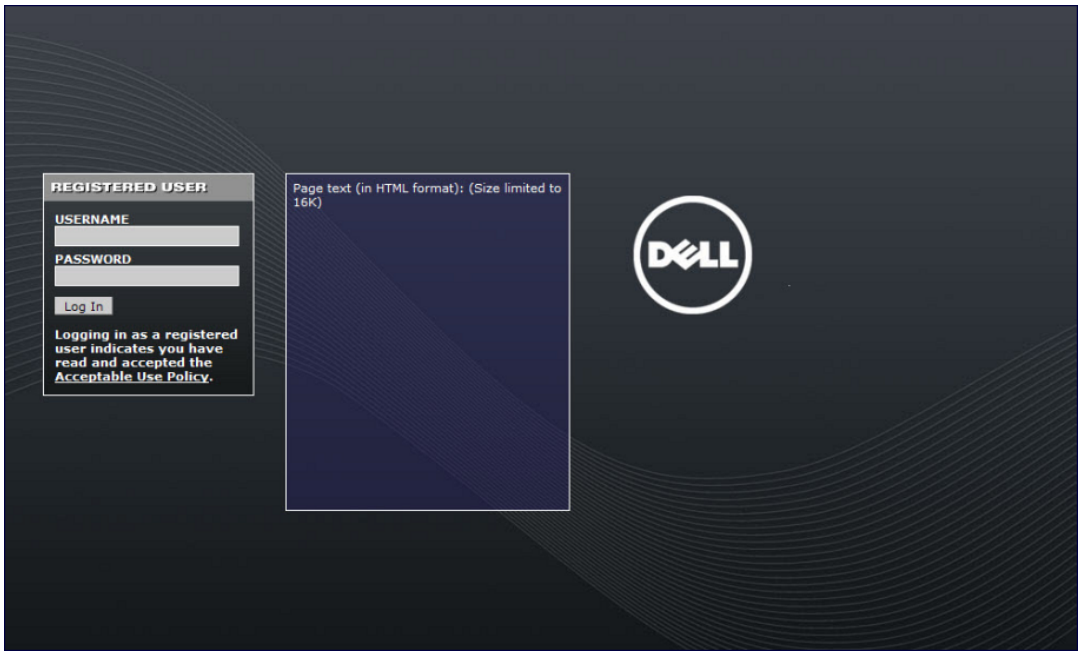
2. To customize the page background:
   a. Select the **YOUR CUSTOM BACKGROUND** page.
   b. Under **Additional options**, enter the location of the JPEG image in the Upload your own custom background field.

c. Set the background color in the Custom page background color field. The color code must a hexadecimal value in the format #hhhhhh.

d. To view the page background changes, click **Submit** at the bottom on the page and then click the **View CaptivePortal** link. The **User Agreement Policy** page appears and displays the Captive Portal page as it will be seen by users.



3. To customize the captive portal background text:
   a. Enter the text that needs to be displayed in the **Page Text (in HTML format)** message box.
   b. To view the background text changes, click **Submit**at the bottom on the page and then click the **View CaptivePortal** link. The **User Agreement Policy** page appears.
   c. Click **Accept.** This displays the Captive Portal page as it will be seen by users.
4. To customize the text under the **Acceptable Use Policy:**
   a. Enter the policy information in the **Policy Text** text box. Use this only in the case of guest logon.
   b. To view the use policy information changes, click **Submit**at the bottom on the page and then click the **View CaptivePortal** link. The **User Agreement Policy** page appears. The text you entered appears in the **Acceptable Use Policy** text box.
   c. Click **Accept.** This displays the Captive Portal page as it will be seen by users



Creating and Installing an Internal Captive Portal

If you do not wish to customize the default captive portal page, you can use the following procedures to create and install a new internal captive portal page. This section describes the following topics:

- Creating a New Internal Web Page on page 323
- Installing a New Captive Portal Page on page 325
- Displaying Authentication Error Messages on page 325
- Reverting to the Default Captive Portal on page 326
- Configuring Localization on page 326
- Customizing the Welcome Page on page 329
- Customizing the Pop-Up box on page 330
- Customizing the Logged Out Box on page 331

Creating a New Internal Web Page

In addition to customizing the default captive portal page, you can also create your own internal web page. A custom web page must include an authentication form to authenticate a user. The authentication form can include any of the following variables listed in Table 54:

Table 54: *Web Page Authentication Variables*

Variable	Description
user	(Required)
password	(Required)
FQDN	The fully-qualified domain name (this is dependent on the setting of the controller and is supported only in Global Catalog Servers software.

The form can use either the "get" or the "post" methods, but the "post" method is recommended. The form's action must absolutely or relatively reference https://<controller_IP>/auth/index.html/u.

You can construct an authentication form using the following HTML:

```
<FORM method="post" ACTION="/auth/index.html/u">
...
</FORM>
```

A recommended option for the <FORM> element is:

```
autocomplete="off"
```

This option prevents Internet Explorer from caching the form inputs. The form variables are input using any form control method available such as INPUT, SELECT, TEXTAREA, and BUTTON. Example HTML code follows.

Username Example

Minimal:

```
<INPUT type="text" name="user">
```

Recommended Options:

```
accesskey="u"   Sets the keyboard shortcut to 'u'
SIZE="25                "Sets the size of the input box to 25
VALUE=          ""Ensures no default value
```

Password Example

Minimal:

```
<INPUT type="password" name="password">
```

Recommended Options:

```
accesskey="p"   Sets the keyboard shortcut to 'p'
SIZE="25                "Sets the size of the input box to 25
VALUE=          ""Ensures no default value
```

FQDN Example

Minimal:

```
<SELECT name=fqdn>
        <OPTION value="fqdn1" SELECTED>
        <OPTION value="fqdn2">
</SELECT>
```

Recommended Options:

```
None
```

Finally, an HTML also requires an input button:

```
<INPUT type="submit">
```

Basic HTML Example

```
<HTML>
  <HEAD>
  </HEAD>
  <BODY>
    <FORM method="post" autocomplete="off" ACTION="/auth/index.html/u">

    Username:<BR>
    <INPUT type="text" name="user" accesskey="u" SIZE="25" VALUE="">
    <BR>

    Password:<BR>
```

```
    <INPUT type="password" name="password" accesskey="p" SIZE="25"
        VALUE="">
    <BR>

    <INPUT type="submit">
    </FORM>
  </BODY>
</HTML>
```

You can find a more advanced example simply by using your browser's "view-source" function while viewing the default captive portal page.

Installing a New Captive Portal Page

You can install the captive portal page by using the Maintenance function of the WebUI.

Log into the WebUI and navigate to **Configuration > Management > Captive Portal > Upload Custom Login Pages**.

This page lets you upload your own files to the controller. There are different page types that you can choose:

- Captive Portal Login (top level): This type uploads the file into the controller and sets the captive portal page to reference the file that you are uploading. Use with caution on a production controller as this takes effect immediately.
- Captive Portal Welcome Page: This type uploads the file that appears after logon and before redirection to the web URL. The display of the welcome page can be disabled or enabled in the captive portal profile.
- Content: The content page type allows you to upload all miscellaneous files that you need to reference from your main captive portal login page. This can be used for images, scripts or any other file that you need to reference. These files are uploaded into the same directory as the top level captive portal page and thus all files can be referenced relatively.

Uploaded files can be referenced using:

```
https://<controller_IP>/upload/custom/<CP-Profile-Name>/<file>
```

Displaying Authentication Error Messages

This section contains a script that performs the following tasks:

- When the user is redirected to the main captive portal login when there is authentication failure, the redirect URL includes a query parameter "errmsg" which java script can extract and display.
- Store the originally requested URL in a cookie so that once the user has authenticated, they are automatically redirected to its original page. Note that for this feature to work, you need ArubaOS release 2.4.2.0 or later. If you don't want this feature, delete the part of the script shown in red.

```
<script>
{
function createCookie(name,value,days)
{
            if (days)
            {
                        var date = new Date();
                        date.setTime(date.getTime()+(days*24*60*60*1000));
                        var expires = "; expires="+date.toGMTString();
            }
            else var expires = "";
            document.cookie = name+"="+value+expires+"; path=/";
}
  var q = window.location.search;
  var errmsg = null;

  if (q && q.length > 1) {
```

```
    q = q.substring(1).split(/[=&]/);
    for (var i = 0; i < q.length - 1; i += 2) {
      if (q[i] == "errmsg") {
        errmsg = unescape(q[i + 1]);
          break;
      }
       if (q[i] == "host") {
          createCookie('url',unescape(q[i+1]),0)
        }
    }
  }

  if (errmsg && errmsg.length > 0) {
    errmsg = "<div id='errorbox'>\n" + errmsg + "\n</div>\n";
    document.write(errmsg);
  }
}
</script>
```

Reverting to the Default Captive Portal

You can reassign the default captive portal site using the "Revert to factory default settings" check box in the "Upload Custom Login Pages" section of the Maintenance tab in the WebUI.

Configuring Localization

The ability to customize the internal captive portal provides you with a very flexible interface to the Dell captive portal system. However, other than posting site-specific messages onto the captive portal website, the most common type of customization is likely to be language localization. This section describes a simple method for creating a native language captive portal implementation using the Dell internal captive portal system.

1. Customize the configurable parts of the captive portal settings to your liking. To do this, navigate to the **Configuration > Management > Captive Portal > Customize Login Page** in the WebUI:

   For example, choose a page design, upload a custom logo and/or a custom background. Also include any page text and acceptable use policy that you would like to include. Put this in your target language or else you will need to translate this at a later time.

   Ensure that Guest login is enabled or disabled as necessary by navigating to the **Configuration > Security > Authentication > L3 Authentication > Captive Portal Authentication Profile** page to create or edit the captive portal profile. Select or deselect "Guest Login".
2. Click **Submit** and then click on **View Captive Portal.** Check that your customization and text/html is correct, with the default interface still in English and the character set still autodetects to ISO-8859-1.

   Repeat steps 1 and 2 until you are satisfied with your page.
3. Once you have a page you find acceptable, click on **View Captive Portal** one more time to display your login page. From your browser, choose "View->Source" or its equivalent. Your system will display the HTML source for the captive portal page. Save this source as a file on your local system.
4. Open the file that you saved in step 3 on page 326, using a standard text editor, and make the following changes:

   a. Fix the character set. The default <HEAD>...</HEAD> section of the file will appear as:

```
<head>
<title>Portal Login</title>

<link href="default1/styles.css" rel="stylesheet" media="screen" type="text/css" />
<script language="javascript" type="text/javascript">
        function showPolicy()
```

```
  {win = window.open("/auth/acceptableusepolicy.html", "policy",
"height=550,width=550,scrollbars=1");}
</script>

</head>
```

In order to control the character set that the browser will use to show the text with, you will need to insert the following line inside the <HEAD>...</HEAD> element:

```
<meta http-equiv="Content-Type" content="text/html; charset=Shift_JIS"/>
```

Replace the "Shift_JIS" part of the above line with the character set that is used by your system. In theory, any character encoding that has been registered with IANA can be used, but you must ensure that any text you enter uses this character set and that your target browsers support the required character set encoding.

b. The final <HEAD>...</HEAD> portion of the document should look similar to this:

```
<head>
<meta http-equiv="Content-Type" content="text/html; charset=Shift_JIS"/>
<title>Portal Login</title>

<link href="default1/styles.css" rel="stylesheet" media="screen" type="text/css" />
<script language="javascript" type="text/javascript">
        function showPolicy()
         {win = window.open("/auth/acceptableusepolicy.html", "policy",
"height=550,width=550,scrollbars=1");}

</script>
</head>
```

c. Fix references: If you have used the built-in preferences, you will need to update the reference for the logo image and the CSS style sheet.

To update the CSS reference, search the text for "<link href" and update the reference to include "/auth/" in front of the reference. The original link should look similar to the following:

```
<link href="default1/styles.css" rel="stylesheet" media="screen" type="text/css" />
```

This should be replaced with a link like the following:

```
<link href="/auth/default1/styles.css" rel="stylesheet" media="screen" type="text/css" />
```

The easiest way to update the image reference is to search for "src" using your text editor and updating the reference to include "/auth/" in front of the image file. The original link should look similar to the following:

```
<img src="default1/logo.gif"/>
```

This should be replaced with a link like this:

```
<img src="/auth/default1/logo.gif"/>
```

d. Insert javascript to handle error cases:

When the controller detects an error situation, it will pass the user's page a variable called "errmsg" with a value of what the error is in English. Currently, only "Authentication Failed" is supported as a valid error message.

To localize the authentication failure message, replace the following text (it is just a few lines below the <body> tag):

```
<div id="errorbox" style="display: none;">
</div>
```

with the script below. You will need to translate the "Authentication Failed" error message into your local language and add it into the script below where it states: localized_msg="...":

```
<script>
{
  var q = window.location.search;
  var errmsg = null;
  if (q && q.length > 1) {
```

```
    q = q.substring(1).split(/[=&]/);
    for (var i = 0; i < q.length - 1; i += 2) {
      if (q[i] == "errmsg") {
        errmsg = unescape(q[i + 1]);
        break;
      }
    }
  }

  if (errmsg && errmsg.length > 0) {
    switch(errmsg) {
    case "Authentication Failed":
        localized_msg="Authentication Failed";
        break;
    default:
      localised_msg=errmsg;
      break;
    }
    errmsg = "<div id='errorbox'>\n" + localised_msg + "\n</div>\n";
    document.write(errmsg);
  };
}
</script>
```

e. Translate the web page text. Once you have made the changes as above, you only need to translate the rest of the text that appears on the page. The exact text that appears will depend on the controller settings when you originally viewed the captive portal. You will need to translate all relevant text such as "REGISTERED USER", "USERNAME", "PASSWORD", the value="" part of the INPUT type="submit" button and all other text. Ensure that the character set you use to translate into is the same as you have selected in part i) above.

Feel free to edit the HTML as you go if you are familiar with HTML.

5. After saving the changes made in step 4 above, upload the file to the controller using the **Configuration > Management > Captive Portal > Upload Custom Login Pages** section of the WebUI.

   Choose the captive portal profile from the drop-down menu. Browse your local computer for the file you saved. For Page Type, select “Captive Portal Login”. Ensure that the "Revert to factory default settings" box is NOT checked and click **Apply**. This will upload the file to the controller and set the captive portal profile to use this page as the redirection page.

   In order to check that your site is operating correctly, go back to the "Customize Login Page" and click on "View Captive Portal" to view the page you have uploaded. Check that your browser has automatically detected the character set and that your text is not garbled.

   To make any adjustments to your page, edit your file locally and simply re-upload to the controller in order to view the page again.

6. Finally, it is possible to customize the welcome page on the controller, however for language localization it is recommended to use an "external welcome page" instead. This can be a web site on an external server, or it can be a static page that is uploaded to a controller.

   You set the welcome page in the captive portal authentication profile. This is the page that the user will be redirected to after successful authentication.

   If this is required to be a page on the controller, the user needs to create their own web page (using the charset meta attribute in step 4 above). Upload this page to the designated controller in the same manner as uploading the captive portal login page under "**Configuration > Management > Captive Portal > Upload Custom Login Pages.** For Page Type, select “Captive Portal Welcome Page”.

Any required client side script (CSS) and media files can also be uploaded using the “Content” Page Type, however file space is limited (use the CLI command **show storage** to see available space). Remember to leave ample room for system files.



The "Registered User" and "Guest User" sections of the login page are implemented as graphics files, referenced by the default CSS styles. In order to change these, you will need to create new graphic files, download the CSS file, edit the reference to the graphics files, change the style reference in your index file and then upload all files as "content" to the controller.

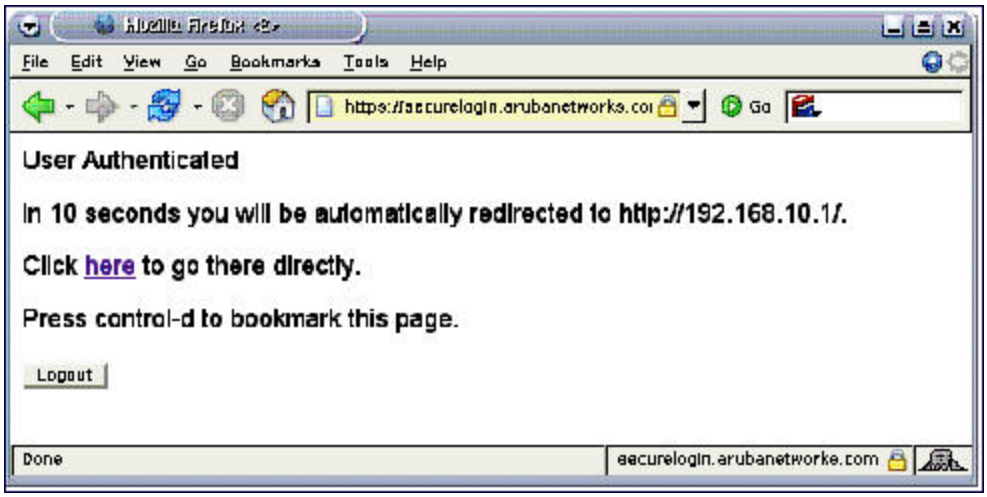
A sample of a translated page is displayed in Figure 34.

Figure 34 *Sample Translated Page*

Customizing the Welcome Page

Once a user is authenticated by the controller, a Welcome page is launched. The default welcome page depends on your configuration, but will look similar to Figure 35:

Figure 35 *Default Welcome Page*



You can customize this welcome page by building your own HTML page and uploading it to the controller. You upload it to the controller by navigating to Management > Captive Portal > Upload Login Pagesand select “Captive Portal Welcome Page” from the Page Type drop-down menu. This file is stored in a directory called "/upload/" on the controller using the file's original name.

In order to actually use this file, you will need to configure the welcome page on the controller. To do this use the CLI command: "aaa captive-portal welcome-page /upload/welc.html" where "welc.html" is the name of the file that you uploaded, or you can change the Welcome page in the captive portal authentication profile in the WebUI.

An example that will create the same page as displayed in Figure 35 is shown below. The part in red will redirect the user to the web page you originally setup. For this to work, please follow the procedure described above in this document.

```
:

<html>
<head>
<script>
{

function readCookie(name)
{
```

```
            var nameEQ = name + "=";
            var ca = document.cookie.split(';');
            for(var i=0;i < ca.length;i++)
            {
                        var c = ca[i];
                        while (c.charAt(0)==' ') c = c.substring(1,c.length);
                        if (c.indexOf(nameEQ) == 0) return c.substring
(nameEQ.length,c.length);
            }
            return null;
}
var cookieval = readCookie('url');
            if (cookieval.length>0) document.write("<meta http-equiv=\"refresh\"
content=\"2;url=http://"+cookieval+"\""+">");
            }
</script>
</head>
<body bgcolor=white text=000000>
<font face="Verdana, Arial, Helvetica, sans-serif" size=+1>
  <b>User Authenticated </b>

<p>In 2 seconds you will be automatically redirected to your original web page</p>
<p> Press control-d to bookmark this page.</p>

<FORM ACTION="/auth/logout.html">
                <INPUT type="submit" name="logout" value="Logout">
</FORM>
</font>
</body>
</html>
```

Customizing the Pop-Up box

In order to customize the Pop-Up box, you must first customize your Welcome page. Once you have customized your welcome page, then you can configure your custom page to use a pop-up box. The default HTML for the pop-up box is:

```
<html>
<body bgcolor=white text=000000>
<font face="Verdana, Arial, Helvetica, sans-serif" size=+1>
 <b>Logout</b></font>
 <p>
  <a href="/auth/logout.html"> Click to Logout </a>
</body>
</html>
```

If you wish your users to be able to logout using this pop-up box, then you must include a reference to /auth/logout.html Once a user accesses this URL then the controller will log them out. It is easiest to simply edit the above HTML to suit your users and then upload the resulting file to the controller using the WebUI under **Configuration > Management > Captive Portal > Upload custom pages** and choose "content" as the page type.

Once you have completed your HTML, then you must get the clients to create the pop-up box once they have logged into the controller. This is done by inserting the following code into your welcome page text and re-uploading the welcome page text to your controller.

Common things to change:

- URL: set the URL to be the name of the pop-up HTML file that you created and uploaded. This should be preceded by "/upload/".
- Width: set w to be the required width of the pop-up box.

- Height: set h to be the required height of the pop-up box.
- Title: set the second parameter in the window.open command to be the title of the pop-up box. Be sure to include the quotes as shown:

```
<script language="JavaScript">
 var url="/upload/popup.html";
 var w=210;
 var h=80;
 var x=window.screen.width - w - 20;
 var y=window.screen.height - h - 60;
 window.open(url, 'logout',
"toolbar=no,location=no,width="+w+",height="+h+",top="+y+",left="+x+",screenX="+x+",screenY
="+y);
</script>
```

Customizing the Logged Out Box

In order to customize the Logged Out box, you must first customize your Welcome page and also your Pop-Up box. To customize the message that occurs after you have logged out then you need to replace the URL that the pop-up box will access in order to log out with your own HTML file.

First you must write the HTML web page that will actually log out the user and will also display page that you wish. An example page is shown below. The key part that must be included is the <iframe>..</iframe> section. This is the part of the HTML that actually does the user logging out. The logout is always performed by the client accessing the /auth/logout.html file on the controller and so it is hidden in the html page here in order to get the client to access this page and for the controller to update its authentication status. If a client does not support the iframe tag, then the text between the <iframe> and the </iframe> is used. This is simply a 0 pixel sized image file that references /auth/logout.html. Either method should allow the client to logout from the controller.

Everything else can be customized.

```
<html>
<body bgcolor=white text=000000>

<iframe src='/auth/logout.html' width=0 height=0 frameborder=0><img src=/auth/logout.html
width=0 height=0></iframe>

<P><font face="Verdana, Arial, Helvetica, sans-serif" size=+1>
You have now logged out.</font></P>

<form> <input type="button" onclick="window.close()" name="close" value="Close
Window"></form>

</body>
</html>
```

After writing your own HTML, then you need to ensure that your customized pop-up box will access your new logged out file. In the pop-up box example above, you simply replace the "/auth/logout.html" with your own file that you upload to the controller. For example, if your customized logout HTML is stored in a file called "loggedout.html" then your "pop-up.html" file should reference it like this:

```
<html>
<body bgcolor=white text=000000>
<font face="Verdana, Arial, Helvetica, sans-serif" size=+1>
 <b>Logout</b></font>
 <p>
  <a href="/upload/loggedout.html"> Click to Logout </a>
</body>
</html>
```

Creating Walled Garden Access

On the Internet, a walled garden typically controls a user's access to web content and services. The walled garden directs the user's navigation within particular areas to allow access to a selection of websites or prevent access to other websites.



The Walled Garden feature can be used with the PEFNG or PEFV licenses.

Walled garden access is needed when an external or internal captive portal is used. A common example could be a hotel environment where unauthenticated users are allowed to navigate to a designated login page (for example, a hotel website) and all its contents.

Users who do not sign up for Internet service can view "allowed" websites (typically hotel property websites). The website names must be DNS-based (not IP address based) and support the option to define wildcards.

Note that the walled garden access feature does not support clients that are configured to use HTTP/HTTPS proxy.

When a user attempts to navigate to other websites not configured in the white list walled garden profile, the user is redirected back to the login page. In addition, the black listed walled garden profile is configured to explicitly block navigation to websites from unauthenticated users.

In the WebUI

1. Navigate to **Advanced Services > Stateful Firewall > Destination.**
2. Click **Add** to add a destination name.
3. Select the controller IP version, IPv4 or IPv6, from the **IP Version**drop-down menu.
4. In the **Destination Name** field, enter a name and click **Add**.
5. Select **name**from the **Rule Type** drop-down menu and add a hostname or wildcard with domain name to which an unauthenticated user is redirected.
6. Click **Apply**.
7. Navigate to **Configuration > Security > Authentication > L3 Authentication.**
8. Select **Captive Portal Authentication Profile.**
9. To allow users to access a domain, enter the destination name that contains the allowed domain names in the **White List** field. This stops unauthenticated users from viewing specific domains such as a hotel website.

   A rule in the white list must explicitly permit a traffic session before it is forwarded to the controller. The last rule in the white list denies everything else.
10. To deny users access to a domain, enter the destination name that contains prohibited domain names in the **Black List** field. This prevents unauthenticated users from viewing specific websites.
11. Click **Apply**.

In the CLI

This example configures a destination named Mywhite-list and adds the domain names, example.com and example.net to that destination. It then adds the destination name Mywhite-list (which contains the allowed domain names example.com and example.net) to the white list.

```
(host)(config)# netdestination "Mywhite-list"
(host)(config)#name example.com
(host)(config)#name example.net

(host) (config) #aaa authentication captive-portal default
(host)(Captive Portal Authentication Profile "default")#white-list Mywhite-list
```

Enabling Captive Portal Enhancements

This release of ArubaOS introduces the following enhancements in Captive Portal:

- Location information such as AP name and AP group name have been included in the Captive Portal redirect URL. The following example shows a Captive Portal redirect URL that contains the AP name and the AP group name:

  https://securelogin.example.com/cgi-bin/login?cmd=login&mac=00:24:d7:ed:84:14&ip=10.15.104.13&essid=example-test-tunnel&apname=ap135&apgroup=example&url=http%3A%2F%2Fwww%2Eespncricinfo%2Ecom%2F

- A new option redirect-url is introduced in the Captive Portal Authentication profile which allows you to redirect the users to a specific URL after the authentication is complete.
- Captive Portal Login URL length has been increased from 256 characters to 2048 characters.
- Support for "?" (question mark) inside the Captive Portal login URL has been added.
- A new field, **description** has been introduced in the **netdestination** and **netdestination6** commands to provide a description about the netdestination up to 128 characters long.
- Support for configuring Whitelist in Captive Portal has been introduced.

The Captive Portal enhancements are available on Tunnel and Split-Tunnel forwarding modes.

Configuring the Redirect-URL

You can configure the Captive Portal redirect URL using the following commands:

```
(host) (config) # aaa authentication captive-portal REDIRECT
(host) (Captive Portal Authentication Profile "REDIRECT") #redirect-url <absolute-URL>
```

Example:

```
(host) (config) # aaa authentication captive-portal REDIRECT
(host) (Captive Portal Authentication Profile "REDIRECT") #redirect-url https://test-login.php
```

Configuring the Login URL

You can configure a Captive Portal login URL up to 2048 characters using the following commands:

```
(host) (config) # aaa authentication captive-portal LOGIN
(host) (Captive Portal Authentication Profile "LOGIN")#login-page
"http://10.17.36.100/login.php?isinit=1&mac=00:11:22:33:44:55&loginURL=https://captiveportal-login.test.aero/auth/index.html&originalURL=&statusURL=&error=&logginIn"
```



You can configure the login URL with "?" (question mark) character in it provided the URL containing the question mark is within the double quotes.

Defining Netdestination Descriptions

You can provide a description (up to 128 characters) for the netdestination using the CLI.

Use the following commands to provide description for an IPv4 netdestination:

```
(host) (config) #netdestination Local-Server
(host) (config-dest) #description "This is a local server for IPv4 client registration"
```

Use the following commands to provide description for an IPv6 netdestination:

```
(host) (config) #netdestination6 Local-Server6
(host) (config-dest) #description "This is a local server for IPv6 client registration"
```

The following command displays the details of the specified IPv4 netdestination:

```
(host) (config-dest) #show netdestination Local-Server

Local-Server Description: This is a local server for IPv4 client registration
------------------------------------------------------------------------------
Position  Type  IP addr   Mask-Len/Range
--------  ----  -------   --------------
1         name  0.0.0.1   yahoomail
2         name  0.0.0.2   mycorp
3         name  0.0.0.3   cricinfo
```

The following command displays the details of the specified IPv6 netdestination:

```
(host) (config-dest) #show netdestination Local-Server6

Local-Server6 Description: This is a local server for IPv6 client registration
------------------------------------------------------------------------------
Position  Type  IP addr   Mask-Len/Range
--------  ----  -------   --------------
1         name  0.0.0.1   yahoomail
2         name  0.0.0.2   mycorp
3         name  0.0.0.3   cricinfo
```

Configuring a Whitelist

You can now configure a Whitelist in Captive Portal using the CLI.

Configuring the Netdestination for a Whitelist:

Use the following commands to configure a netdestination alias for Whitelist:

```
(host) (config) #netdestination whitelist
(host) (config-dest) #description guest_whitelist
(host) (config-dest) #name mycorp
```

Associating a Whitelist to Captive Portal Profile

Use the following CLI commands to associate a whitelist to the Captive profile:

```
(host) (config) #aaa authentication captive-portal CP_Profile
(host) (Captive Portal Authentication Profile "CP_Profile”) #white-list whitelist
```

Applying a Captive Portal Profile to a User-Role

Use the following commands to apply the Captive Portal profile to a user-role:

```
(host) (config) # user-role guest_role
(host) (config-role) #session-acl logon-control
(host) (config-role) #session-acl captiveportal
(host) (config-role) #captive-portal CP_Profile
```

Verifying a Whitelist Configuration

Use the following commands to verify the whitelist alias:

```
(host) (config) #show netdestination whitelist

whitelist Description: guest_whitelist
--------------------------------------
Position  Type  IP addr   Mask-Len/Range
--------  ----  -------   --------------
1         name  0.0.0.6   mycorp
```

Verifying a Captive Portal Profile Linked to a Whitelist

Use the following commands to verify the Captive Portal profile linked to the whitelist:

```
(host) (config) #show aaa authentication captive-portal CP_Profile

Captive Portal Authentication Profile "CP_Profile"
------------------------------------------------------------------
Parameter                                          Value
---------                                          -----
Default Role                                       guest
Default Guest Role                                 guest
Server Group                                       default
Redirect Pause                                     10 sec
User Login                                         Enabled
Guest Login                                        Disabled
Logout popup window                                Enabled
Use HTTP for authentication                        Disabled
Logon wait minimum wait                            5 sec
Logon wait maximum wait                            10 sec
logon wait CPU utilization threshold               60 %
Max Authentication failures                        0
Show FQDN                                          Disabled
Use CHAP (non-standard)                            Disabled
Login page                                         /auth/index.html
Welcome page                                       /auth/welcome.html
Show Welcome Page                                  Yes
Add switch IP address in the redirection URL       Disabled
Adding user vlan in redirection URL                Disabled
Add a controller interface in the redirection URL  N/A
Allow only one active user session                 Disabled
White List                                         whitelist
Black List                                         N/A
Show the acceptable use policy page                Disabled
Redirect URL                                       N/A
```

Verifying Dynamic ACLs for a Whitelist

Use the following commands to verify the dynamically created ACLs for the whitelist:

```
(host) (config) #show rights guest_role

Derived Role = 'guest_role'
Up BW:No Limit   Down BW:No Limit
L2TP Pool = default-l2tp-pool
PPTP Pool = default-pptp-pool
Periodic reauthentication: Disabled
ACL Number = 79/0
Max Sessions = 65535
Captive Portal profile = CP_Profile

access-list List
----------------
Position  Name                                    Location
--------  ----                                    --------
1         CP_Profile_list_operations
2         logon-control
3         captiveportal
CP_Profile_list_operations
----------------------------------------
Priority  Source  Destination  Service    Action  TimeRange  Log  Expired  Queue  TOS  8021P
Blacklist  Mirror  DisScan  ClassifyMedia  IPv4/6
--------  ------  -----------  -------    ------  ---------  ---  -------  -----  ---  -----
---------  ------  -------  -------------  ------
1         user    whitelist    svc-http   permit                          Low
                                          4
```

```
2         user    whitelist    svc-https  permit                          Low
                                          4

logon-control
-------------
Priority  Source  Destination  Service   Action  TimeRange  Log  Expired  Queue  TOS  8021P
Blacklist  Mirror  DisScan  ClassifyMedia  IPv4/6
--------  ------  -----------  -------   ------  ---------  ---  -------  -----  ---  -----  -
--------  ------  -------  -------------  ------
1         user    any          udp 68    deny                            Low
                                          4
2         any     any          svc-icmp  permit                          Low
                                          4
3         any     any          svc-dns   permit                          Low
                                          4
4         any     any          svc-dhcp  permit                          Low
                                          4
5         any     any          svc-natt  permit                          Low
                                          4

captiveportal
-------------
Priority  Source  Destination  Service          Action        TimeRange  Log  Expired  Queue
TOS  8021P  Blacklist  Mirror  DisScan  ClassifyMedia  IPv4/6
--------  ------  -----------  -------          ------        ---------  ---  -------  -----
---  -----  ---------  ------  -------  -------------  ------
1         user    controller   svc-https        dst-nat 8081                        Low
                                                      4
2         user    any          svc-http         dst-nat 8080                        Low
                                                      4
3         user    any          svc-https        dst-nat 8081                        Low
                                                      4
4         user    any          svc-http-proxy1  dst-nat 8088                        Low
                                                      4
5         user    any          svc-http-proxy2  dst-nat 8088                        Low
                                                      4
6         user    any          svc-http-proxy3  dst-nat 8088                        Low
                                                      4
Expired Policies (due to time constraints) = 0
```

Verifying DNS Resolved IP Addresses for Whitelisted URLs

Use the following command to verify the DNS resolved IP addresses for the whitelisted URLs:

```
(host) #show firewall dns-names ap-name <AP-name>
```

Example:

```
(host) #show firewall dns-names ap-name ap135

Firewall DNS names
------------------
Index          Name                 Id           Num-IP       List
-----          ----                 --           ------       ----
0              bugzilla             10                1    0.0.0.0
1              cricinfo              9                0
2              yahoo                 1                0
3              mycorp                6                1    1.1.1.1
```

Wireless networks can use virtual private network (VPN) connections to further secure wireless data from attackers. The Dell controller can be used as a VPN concentrator that terminates all VPN connections from both wired and wireless clients.

This chapter describes the following topics:

- Planning a VPN Configuration on page 337
- Working with VPN Authentication Profiles on page 340
- Configuring a Basic VPN for L2TP/IPsec in the WebUI on page 342
- Configuring a VPN for L2TP/IPsec with IKEv2 in the WebUI on page 346
- Configuring a VPN for Smart Card Clients on page 350
- Configuring a VPN for Clients with User Passwords on page 351
- Configuring Remote Access VPNs for XAuth on page 353
- Working with Remote Access VPNs for PPTP on page 355
- Working with Site-to-Site VPNs on page 355
- Working with VPN Dialer on page 361

Planning a VPN Configuration

You can configure the controller for the following types of VPNs:

- Remote access VPNs allow hosts (for example, telecommuters or traveling employees) to connect to private networks (for example, a corporate network) over the Internet. Each host must run VPN client software which encapsulates and encrypts traffic, then sends it to a VPN gateway at the destination network. The controller supports the following remote access VPN protocols:
  - Layer-2 Tunneling Protocol over IPsec (L2TP/IPsec)
  - Point-to-Point Tunneling Protocol (PPTP)
  - XAUTH IKE/IPsec
  - IKEv2 with Certificates
  - IKEv2 with EAP
- Site-to-site VPNs allow networks (for example, a branch office network) to connect to other networks (for example, a corporate network). Unlike a remote access VPN, hosts in a site-to-site VPN do not run VPN client software. All traffic for the other network is sent and received through a VPN gateway, which encapsulates and encrypts the traffic.

Before enabling VPN authentication, you must configure the following:

- The default user role for authenticated VPN clients. See Roles and Policies on page 364 for information about configuring user roles.
- The authentication server group the controller uses to validate the clients. See Authentication Servers on page 225 for configuration details.



A server-derived role, if present, takes precedence over the default user role.

You then specify the default user role and authentication server group in the VPN authentication **default** profile, as described in the sections below.



ESP Tunnel Mode is the only supported IPsec mode of operation. ArubaOS does not support AH and Transport modes.

Selecting an IKE protocol

Controllers running ArubaOS version 6.1 and later support both IKEv1 and the newer IKEv2 protocol to establish IPsec tunnels. IKEv2 is simpler, faster, and a more reliable protocol than IKEv1, though both IKEv1 and IKEv2 support the same suite-B cryptographic algorithms.

If your IKE policy uses IKEv2, you should be aware of the following caveats when you configure your VPN:

- ArubaOS does not support separate pre-shared keys for both directions of an exchange; both peers must use the same pre-shared key. ArubaOS does not support mixed authentication with both pre-shared keys and certificates; each authentication exchange requires a single authentication type. For example, if a client authenticates with a pre-shared key, the controller must also authenticate with a pre-shared key.
- ArubaOS does not support IKEv2 mobility (MOBIKE), Authentication Headers (AH), or IP Payload Compression Protocol (IPComp).

Understanding Suite-B Encryption Licensing

Dell controllers support Suite-B cryptographic algorithms when the Advanced Cryptography (ACR) license is installed. Table 55 describes the Suite-B algorithms supported by ArubaOS IKE Policies and IPsec tunnels. For further details on configuring a VPN to use Suite-B algorithms, see Configuring a VPN for L2TP/IPsec with IKEv2 in the WebUI on page 346.

Table 55: *Suite-B Algorithms Supported by the ACR License*

IKE Policies	Suite-B for IPsec tunnels
hash: SHA-256-128, SHA-384-192	Encryption: AES-128-GCM, AES-256-GCM
Diffie-Hellman (DH) Groups: ECP-256, ECP-384	Perfect Forward Secrecy (PFS): ECP-256, ECP-384
Pseudo-Random Function (PRF): HMAC_SHA_256, HMAC_SHA_384	—
Suite-B certificates: ECDSA-256, ECDSA-384	—



The ArubaOS hardware supports IKE Suite-B AES-128-GCM and AES-256-GCM encryption. ArubaOS software performs the KE Suite-B Diffie-Hellman and Certificate-based signature operations, and hash, PFS, and PRF algorithm functions.

The following VPN clients support Suite-B algorithms when establishing an L2TP/IPsec VPN:

Table 56: *Client Support for Suite-B*

Client Operating System	Supported Suite-B IKE Authentication	Supported Suite-B IPsec Encryption
• Windows 7 • Windows Vista • Windows XP	• IKEv1 Clients using ECDSA Certificates • IKEv1/IKEv2 Clients using ECDSA Certificates with L2TP/PPP/EAP-TLS certificate user-authentication	• AES-128-GCM • AES-256-GCM

The Suite-B algorithms described in Table 55 are also supported by Site-to-Site VPNs between Dell controllers, or between a Dell controller and a server running Windows 2008 or StrongSwan 4.3.

Working with IKEv2 Clients

Not all clients support the both the IKEv1 and IKEv2 protocols. Only the clients in Table 57 support IKEv2 with the following authentication types:

Table 57: *VPN Clients Supporting IKEv2*

Windows 7 Client	StrongSwan 4.3 Client	VIA Client
• Machine authentication with Certificates • User name password authentication using EAP-MSCHAPv2 or PEAP-MSCHAPv2 • User smart-card authentication with EAP-TLS / IKEv2 **NOTE:** Windows 7 clients using IKEv2 do not support pre-shared key authentication.	• Machine authentication with Certificates • User name password authentication using EAP-MSCHAPv2 • Suite-B cryptographic algorithms	• Machine authentication with Certificates • User name password authentication using EAP-MSCHAPv2 • EAP-TLS using Microsoft cert repository **NOTE:** VIA clients using IKEv2 do not support pre-shared key authentication.

Understanding Supported VPN AAA Deployments

If you want to simultaneously deploy various combinations of a VPN client, RAP-psk, RAP-certs, and CAP on the same controller, see Table 58.

Each row in this table specifies the allowed combinations of AAA servers for simultaneous deployment. Configuration rules include:

- RAP-certs can only use LocalDB-AP.
- A RAP-psk and RAP-cert can only terminate on the same controller if the RAP VPN profile's AAA server uses Local-db.
- If a RAP-psk is using an external AAA server, then the RAP-cert cannot be terminated on the same controller.
- Clients can use any type of AAA server, regardless of the RAP/CAP authentication configuration server.

Table 58: *Supported VPN AAA Deployments*

VPN Client	RAP psk	RAP certs	CAP
External AAA server 1	LocalDB	LocalDB-AP	CPSEC-whitelist
External AAA server 1	External AAA server 1	Not supported	CPSEC-whitelist
External AAA server 1	External AAA server 2	Not supported	CPSEC-whitelist
LocalDB	LocalDB	LocalDB-AP	CPSEC-whitelist
LocalDB	External AAA server 1	Not supported	CPSEC-whitelist

Working with Certificate Groups

The certificate group feature allows you to access multiple types of certificates on the same controller. To create a certificate group, use the following command:

```
(host) (config) #crypto-local isakmp certificate-group server-certificate server_certificate
ca-certificate ca_certificate
```

You can view existing certificate groups using:

```
show crypto-local isakmp certificate-group
```

Working with VPN Authentication Profiles

VPN Authentication profiles identify a user role for authenticated VPN clients, an authentication server, and the server group to which the authentication server belongs. There are three predefined VPN authentication profiles: **default**, **default-rap** and **default-cap**. These different profiles allow you to use different authentication servers, user roles and IP pools for VPN, remote AP, and campus AP clients.



You can configure the **default** and **default-rap** profiles, but not the **default-cap** profile.

Table 59: *Predefined Authentication Profile settings*

Parameter	Description	default	default-rap	default-cap
Default Role for authenticated users	The role that will be assigned to the authenticated users.	default-vpn-role	default-vpn-role	sys-ap-role 0
Maximum allowed authentication failures	The number of contiguous authentication failures before the station is blacklisted.	0 (feature is disabled)	0 (feature is disabled)	0 (feature is disabled)
Check certificate common name against AAA server		disabled	enabled	enabled
Export VPN IP address as a route	When enabled, this causes any VPN client address to be exported to	enabled	enabled	enabled

Parameter	Description	default	default-rap	default-cap
	OSPF using IPC. **NOTE:** Note that the Framed-IP-Address attribute is assigned the IP address as long as the any server returns the attribute. The Framed-IP-Address value always has a higher priority than the local address pool.			
User idle timeout	The user idle timeout value for this profile. Specify the idle timeout value for the client in seconds. Valid range is 30-15300 in multiples of 30 seconds. Enabling this option overrides the global settings configured in the AAA timers. If this is disabled, the global settings are used.	disabled	N/A	N/A
PAN firewalls Integration	Requires IP mapping at Palo Alto Networks firewalls.	disabled	disabled	disabled

To edit the **default** VPN authentication profile:

1. Navigate to the **Configuration > Advanced Services > All Profiles > VPN Authentication > default**page.
2. In the **Profiles** list in the left window pane, select the **default** VPN Authentication Profile.
3. Click the **Default Role** drop-down list, and select the default user role for authenticated VPN users. (For detailed information on creating and managing user roles and policies, see Roles and Policies on page 364.)
4. (Optional) If you use client certificates for user authentication, select the **Check certificate common name against AAA server** checkbox to verify that the certificate's common name exists in the server. This parameter is enabled by default in the **default-cap** and **default-rap** VPN profiles, and disabled by default on all other VPN profiles.
5. (Optional) Set **Max Authentication failures** to an integer value. The default value is 0, which disables this feature.
6. (Optional) Regardless of how an authentication server is contacted, the **Export VPN IP address as a route** option causes any VPN client address to be exported to OSPF using IPC. Note that the Framed-IP-Address attribute is assigned the IP address as long as any server returns the attribute. The Framed-IP-Address value always has a higher priority than the local address pool.
7. (Optional) Enabling PAN firewalls Integration requires IP mapping at Palo Alto Networks firewalls. (For more information about PAN firewall integration, see Palo Alto Networks Firewall Integration on page 620.)
8. Click **Apply**.
9. In the **Default** profile menu in the left window pane, select **Server Group**.

10. From the **Server Group** drop-down list, select the server group to be used for VPN authentication.
11. Click **Apply**.

To configure VPN authentication via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #aaa authentication vpn default
   cert-cn-lookup
   clone
   default-role <role>
   export-route
   max-authentication-failure <number>
   pan-integration
   radius-accounting <server_group_name>
   server-group <name>
   user-idle-timeout <seconds>
```

Configuring a Basic VPN for L2TP/IPsec in the WebUI

The combination of Layer-2 Tunneling Protocol and Internet Protocol Security (L2TP/IPsec) is a highly-secure technology that enables VPN connections across public networks such as the Internet. L2TP/IPsec provides both a logical transport mechanism on which to transmit PPP frames, tunneling, or encapsulation, so that the PPP frames can be sent across an IP network. L2TP/IPsec relies on the PPP connection process to perform user authentication and protocol configuration. With L2TP/IPsec, the user authentication process is encrypted using the Data Encryption Standard (DES) or Triple DES (3DES) algorithm.

L2TP/IPsec using IKEv1 requires two levels of authentication:

- Computer-level authentication with a preshared key to create the IPsec security associations (SAs) to protect the L2TP-encapsulated data.
- User-level authentication through a PPP-based authentication protocol using passwords, SecureID, digital certificates, or smart cards after successful creation of the SAs.



Note that only Windows 7 clients, StrongSwan 4.3 clients, and VIA clients support IKEv2. For additional information on the authentication types supported by these clients, see Working with IKEv2 Clients on page 339.

Use the following procedures to configure a remote access VPN for L2TP IPsec for clients using pre-shared keys, certificates or EAP for authentication using the WebUI:

- Defining Authentication Method and Server Addresses on page 347
- Defining Address Pools on page 347
- Enabling Source NAT on page 347
- Selecting Certificates on page 347
- Defining IKEv1 Shared Keys on page 344
- Configuring IKE Policies on page 348
- Setting the IPsec Dynamic Map on page 349
- Finalizing WebUI changes on page 350

Defining Authentication Method and Server Addresses

1. Define the authentication method and server addresses.
2. Navigate to **Configuration > Advanced Services > VPN Services** and click the **IPSEC** tab.
3. To enable L2TP, select **Enable L2TP** (this is enabled by default).
4. Select the authentication method for IKEv1 clients. Currently supported methods are:

- Password Authentication Protocol (PAP)
- Extensible Authentication Protocol (EAP)
- Challenge Handshake Authentication Protocol (CHAP)
- Microsoft Challenge Handshake Authentication Protocol (MSCHAP)

5. Configure the IP addresses of the primary and secondary Domain Name System (DNS) servers and the primary and secondary Windows Internet Naming Service (WINS) Server that are pushed to the VPN client.

Defining Address Pools

Next, define the pool from which the clients are assigned addresses.

1. In the **Address Pools** section of the **IPSEC** tab, click **Add** to open the **Add Address Pool** page.
2. Specify the pool name, the start address, and the end address.
3. Click **Done**.

RADIUS Framed-IP-Address for VPN Clients

IP addresses are usually assigned to VPN clients from configured local address pools. However, the Framed-IP-Address attribute that is returned from a RADIUS server can be used to assign the address.

VPN clients use different mechanisms to establish VPN connections with the controller, such as IKEv1, IKEv2, EAP, or a user certificate. Regardless of how the RADIUS server is contacted for authentication, the Framed-IP-Address attribute is assigned the IP address as long as the RADIUS server returns the attribute. The Framed-IP-Address value always has a higher priority than the local address pool.

Enabling Source NAT

In the **Source NAT** section of the **IPSEC** tab, select **Enable Source NAT** if the IP addresses of clients need to be translated to access the network. If you enabled source NAT, click the **NAT pool** drop-down list and select an existing NAT pool. If you have not yet created the NAT pool you want to use:

1. Navigate to **Configuration > IP > NAT Pools**.
2. Click **Add**.
3. In the **Pool Name** field, enter a name for the new NAT pool, up to 63 alphanumeric characters.
4. In the **Start IP address** field, enter the dotted-decimal IP address that defines the beginning of the range of source NAT addresses in the pool.
5. In the **End IP address** field, enter the dotted-decimal IP address that defines the end of the range of source NAT addresses in the pool.
6. In the **Destination NAT IP Address** field, enter the destination NAT IP address in dotted-decimal format. If you do not enter an address into this field, the NAT pool uses the destination NAT IP **0.0.0.0**.
7. Click **Done.**
8. Navigate to **Configuration > Advanced Services > VPN Services** and click the **IPSEC** tab to return to the **IPsec** window.
9. Click the **NAT Pool** drop-down list and select the NAT pool you just created.

Selecting Certificates

If you are configuring a VPN to support machine authentication using certificates, define the IKE Server certificates for VPN clients using IKE. Note that these certificate must be imported into the controller, as described in Management Access on page 778.

1. Select the server certificate for client machines using IKE by clicking **the IKE Server Certificate** drop-down list and selecting an available certificate name.

2. If you are configuring a VPN to support clients using certificates, you must also assign one or more trusted CA certificates to VPN clients.
   a. Under CA Certificate Assigned for VPN-clients, click **Add**.
   b. Select a CA certificate from the drop-down list of CA certificates imported in the controller.
   c. Click **Done**.
   d. Repeat the above steps to add additional CA certificates.

Defining IKEv1 Shared Keys

If you are configuring a VPN to support IKEv1 and clients using pre-shared keys, You can configure a global IKE key or configure an IKE key for each subnet. Make sure that this key matches the key on the client.

1. In the **IKE Shared Secrets** section of the **IPsec** tab, click **Add** to open the **Add IKE Secret page**.
2. Enter the subnet and subnet mask. To make the IKE key global, specify 0.0.0.0 for both values.
3. Enter the IKE Shared Secret and Verify IKE Shared Secret.
4. Click **Done**.

Configuring IKE Policies

ArubaOS contains several predefined default IKE policies, as described in Table 60. If you do not want to use any of these predefined policies, you can use the procedures below to edit an existing policy or create your own custom IKE policy instead.



The IKE policy selections, along with any preshared key, need to be reflected in the VPN client configuration. When using a third-party VPN client, set the VPN configuration on clients to match the choices made above. In case the Dell dialer is used, these configurations need to be made on the dialer prior to downloading the dialer onto the local client.

1. Scroll down to the **IKE Policies** section of the **IPSEC** tab, then click **Edit** to edit an existing policy or click **Add** to create a new policy.
2. Enter a number into the **Priority** field to set the priority for this policy. Enter a priority of 1 for the configuration to take priority over the Default setting.
3. Select the IKE version. Click the **Version** drop-down list and select **V1** for IKEv1 or **V2** for IKEv2.
4. Set the Encryption type. Click the **Encryption** drop-down list and select one of the following encryption types:
   - DES
   - 3DES
   - AES128
   - AES192
   - AES256
5. Set the HASH function. Click the **Hash** drop-down list and select one of the following hash types:
   - MD5
   - SHA
   - SHA1-96
   - SHA2-256-128
   - SHA2-384-192
6. ArubaOS VPNs support client authentication using pre-shared keys, RSA digital certificates, or Elliptic Curve Digital Signature Algorithm (ECDSA) certificates. To set the authentication type for the IKE rule, click the **Authentication** drop-down list and select one of the following types:

- Pre-Share (for IKEv1 clients using pre-shared keys)
- RSA (for clients using certificates)
- ECDSA-256 (for clients using certificates)
- ECDSA-384 (for clients using certificates)

7. Diffie-Hellman is a key agreement algorithm that allows two parties to agree upon a shared secret, and is used within IKE to securely establish session keys. To set the Diffie–Hellman Group for the ISAKMP policy, click the **Diffi–Hellman Group** drop-down list and select one of the following groups:
   - Group 1: 768-bit Diffie–Hellman prime modulus group.
   - Group 2: 1024-bit Diffie–Hellman prime modulus group.
   - Group 14: 2048-bit Diffie–Hellman prime modulus group.
   - Group 19: 256-bit random Diffie–Hellman ECP modulus group.
   - Group 20: 384-bit random Diffie–Hellman ECP modulus group.
8. Set the **Security Association Lifetime** to define the lifetime of the security association, in seconds. The default value is 7200 seconds. To change this value, uncheck the **default** checkbox and enter a value from **300** to **86400** seconds.
9. Click **Done**.

Setting the IPsec Dynamic Map

Dynamic maps enable IPsec SA negotiations from dynamically addressed IPsec peers. ArubaOS has a predefined IPsec dynamic map for IKEv1. If you do not want to use this predefined map, you can use the procedures below to edit an existing map, or create your own custom IPsec dynamic map instead.

1. Scroll down to the IPsec Dynamic Map section of the IPSEC tab, then click **Edit** by a map name to edit an existing map or click **Add** to create a new map.
2. In the **Name** field, enter a name for the dynamic map.
3. In the **Priority** field, enter a priority number for the map. Negotiation requests for security associations try to match the highest-priority map first. If that map does not match, the negotiation request continues down the list to the next-highest priority map until a match is made.
4. Click the **Version** drop-down list and select **V1** to create an IPsec map for remote peers using IKEv1.
5. (Optional) Configure Perfect Forward Secrecy (PFS) settings for the dynamic peer by assigning a Diffie-Hellman prime modulus group. PFS provides an additional level of security by ensuring that the IPsec SA key was not derived from any other key, and therefore can not be compromised if another key is broken. Click the **Set PFS** drop-down list and select one of the following groups:
   - Group 1: 768-bit Diffie–Hellman prime modulus group.
   - Group 2: 1024-bit Diffie–Hellman prime modulus group.
   - Group 14: 2048-bit Diffie–Hellman prime modulus group.
   - Group 19: 256-bit random Diffie–Hellman ECP modulus group.
   - Group 20: 384-bit random Diffie–Hellman ECP modulus group.
6. Select the transform set for the map to define a specific encryption and authentication type used by the dynamic peer. Click the **Transform Set** drop-down list, and select the transform set for the dynamic peer.



To view current configuration settings for an IPsec transform-set, access the command-line interface and issue the command `crypto ipsec transform-set tag <transform-set-name>`.

7. Set the **Security Association Lifetime** to define the lifetime of the security association for the dynamic peer, in seconds. The default value is 7200 seconds. To change this value, uncheck the **default** checkbox and enter a value from **300** to **86400** seconds.

8. Click **Done**.

Finalizing WebUI changes

When you have finished configuring your IPsec VPN settings, click **Apply** to apply the new settings before navigating to other pages.

Configuring a Basic L2TP VPN in the CLI

Use the following procedures to use the command-line interface to configure a remote access VPN for L2TP IPsec:

1. Define the authentication method and server addresses:

```
(host)(config) #vpdn group l2tp
  enable
  client configuration {dns|wins} <ipaddr1> [<ipaddr2>]
```

2. Enable authentication methods for IKEv1 clients:

```
 vpdn group l2tp ppp authentication {cache-securid|chap|eap|mschap|mschapv2|pap
```

3. Create address pools:

```
(host)(config) #ip local pool <pool> <start-ipaddr> <end-ipaddr>
```

4. Configure source NAT:

```
(host)(config) #ip access-list session srcnatuser any any src-nat pool <pool> position 1
```

5. If you are configuring a VPN to support machine authentication using certificates, define server certificates for VPN clients using IKEv1:

```
(host)(config) #crypto-local isakmp server-certificate <cert>
```

6. If you are configuring a VPN to support IKEv1 Clients using pre-shared keys, you can configure a global IKE key by entering **0.0.0.0** for both the address and netmask parameters in the command below, or configure an IKE key for an individual subnet by specifying the IP address and netmask for that subnet:

```
crypto isakmp key <key> address <ipaddr|> netmask <mask>
```

7. Define IKE Policies:

```
(host)(config) #crypto isakmp policy <priority>
encryption {3des|aes128|aes192|aes256|des}
version v1|v2
authentication {pre-share|rsa-sig|ecdsa-256ecdsa-384}
group {1|2|19|20}
hash {md5|sha|sha1-96|sha2-256-128|sha2-384-192}
lifetime <seconds>
```

Configuring a VPN for L2TP/IPsec with IKEv2 in the WebUI

Only clients running Windows 7, StrongSwan 4.3, and Dell VIA support IKEv2. For additional information on the authentication types supported by these clients, see "Working with IKEv2 Clients on page 339."

Use the following procedures to in the WebUI configure a remote access VPN for IKEv2 clients using certificates.

- Defining Authentication Method and Server Addresses on page 347
- Defining Address Pools on page 347
- Enabling Source NAT on page 347
- Selecting Certificates on page 347
- Configuring IKE Policies on page 348
- Setting the IPsec Dynamic Map on page 349
- Finalizing WebUI changes on page 350

Defining Authentication Method and Server Addresses

1. Define the authentication method and server addresses.
2. Navigate to **Configuration > Advanced Services > VPN Services** and click the **IPSEC** tab.
3. To enable L2TP, select **Enable L2TP** (this is enabled by default).
4. Select the authentication method for IKEv1 clients. Currently supported methods are:
   - Password Authentication Protocol (PAP)
   - Extensible Authentication Protocol (EAP)
   - Challenge Handshake Authentication Protocol (CHAP)
   - Microsoft Challenge Handshake Authentication Protocol (MSCHAP)
   - Microsoft Challenge Handshake Authentication Protocol version 2 (MSCHAPv2)
5. Configure the IP addresses of the primary and secondary Domain Name System (DNS) servers and primary and secondary Windows Internet Naming Service (WINS) Server that is pushed to the VPN client.

Defining Address Pools

Next, define the pool from which the clients are assigned addresses.

1. In the **Address Pools** section of the **IPSEC** tab, click **Add** to open the **Add Address Pool** page.
2. Specify the pool name, the start address, and the end address.
3. Click **Done**.

Enabling Source NAT

In the **Source NAT** section of the **IPSEC** tab, select **Enable Source NAT** if the IP addresses of clients need to be translated to access the network. If you enabled source NAT, click the **NAT pool** drop-down list and select an existing NAT pool. If you have not yet created the NAT pool you want to use:

1. Navigate to **Configuration > IP > NAT Pools**.
2. Click **Add**.
3. In the **Pool Name** field, enter a name for the new NAT pool, up to 63 alphanumeric characters.
4. In the **Start IP address** field, enter the dotted-decimal IP address that defines the beginning of the range of source NAT addresses in the pool.
5. In the **End IP address** field, enter the dotted-decimal IP address that defines the end of the range of source NAT addresses in the pool.
6. In the **Destination NAT IP Address** field, enter the destination NAT IP address in dotted-decimal format. If you do not enter an address into this field, the NAT pool uses the destination NAT IP **0.0.0.0**.
7. Click **Done** to close the NAT pools tab.
8. Navigate to **Configuration > Advanced Services > VPN Services** and click the **IPSEC** tab to return to the **IPSEC** window.
9. Click the **NAT Pool** drop-down list and select the NAT pool you just created.

Selecting Certificates

To configure the VPN to support machine authentication using certificates, define the IKE Server certificates for VPN clients using IKEv2. Note that these certificate must be imported into the controller, as described in Management Access on page 778.

1. Select the IKEv2 server certificate for client machines using IKEv2 by clicking the **IKEv2 Server Certificate** drop-down list and selecting an available certificate name.
2. If you are configuring a VPN to support IKEv2 clients using certificates, you must also assign one or more trusted CA certificates to VPN clients.

a. Under CA Certificate Assigned for VPN-clients, click **Add**.
b. Select a CA certificate from the drop-down list of CA certificates imported in the controller.
c. Click **Done**.
d. Repeat the above steps to add additional CA certificates.

Configuring IKE Policies

ArubaOS contains several predefined default IKE policies, as described in Table 60. If you do not want to use any of these predefined policies, you can use the procedures below to delete a factory-default policy, edit an existing policy, or create your own custom IKE policy instead.



The IKE policy selections need to be reflected in the VPN client configuration. When using a third-party VPN client, set the VPN configuration on clients to match the choices made above. In case the Dell dialer is used, these configurations need to be made on the dialer prior to downloading the dialer onto the local client.

1. Scroll down to the **IKE Policies** section of the **IPSEC** tab, then click **Edit** to edit an existing policy or click **Add** to create a new policy.

   You can also delete a predefined factory-default IKE policy by clicking **Delete**.
2. Enter a number into the **Priority** field to set the priority for this policy. Enter a priority of 1 for the configuration to take priority over the Default setting.
3. Select the IKE version. Click the **Version** drop-down list and select **V2** for IKEv2.
4. Set the Encryption type. Click the **Encryption** drop-down list and select one of the following encryption types:
   - DES
   - 3DES
   - AES128
   - AES192
   - AES256
5. Set the HASH function. Click the **Hash** drop-down list and select one of the following hash types:
   - MD5
   - SHA
   - SHA1-96
   - SHA2-256-128
   - SHA2-384-192
6. ArubaOS VPNs support IKEv2 client authentication using RSA digital certificates, or Elliptic Curve Digital Signature Algorithm (ECDSA) certificates. To set the authentication type for the IKE rule, click the **Authentication** drop-down list and select one of the following types:
   - RSA
   - ECDSA-256
   - ECDSA-384
7. Diffie-Hellman is a key agreement algorithm that allows two parties to agree upon a shared secret, and is used within IKE to securely establish session keys. To set the Diffie–Hellman Group for the ISAKMP policy, click the **Diffie–Hellman Group** drop-down list and select one of the following groups:
   - Group 1: 768-bit Diffie–Hellman prime modulus group.
   - Group 2: 1024-bit Diffie–Hellman prime modulus group.
   - Group 19: 256-bit random Diffie–Hellman ECP modulus group.
   - Group 20: 384-bit random Diffie–Hellman ECP modulus group.

8. Set the Pseudo-Random Function (PRF) value. This algorithm is an HMAC function to used to hash certain values during the key exchange:
   - PRF-HMAC-MD5
   - PRF-HMAC-SHA1
   - PRF-HMAC-SHA256
   - PRF-HMAC-SHA384
9. Set the **Security Association Lifetime** to define the lifetime of the security association, in seconds. The default value is 7200 seconds. To change this value, uncheck the default checkbox and enter a value from **300** to **86400** seconds.
10. Click **Done.**

Setting the IPsec Dynamic Map

Dynamic maps enable IPsec SA negotiations from dynamically addressed IPsec peers. ArubaOS has a predefined IPsec dynamic maps for IKEv2. If you do not want to use of these predefined maps, you can use the procedures below to to delete a factory-default map,edit an existing map or create your own custom IPsec dynamic map instead.

In the WebUI

1. Scroll down to the IPsec Dynamic Map section of the IPSEC tab, then click **Edit** by a map name to edit an existing map or click **Add** to create a new map.



You can also delete a predefined factory-default dynamic map by clicking **Delete**.

2. In the **Name** field, enter a name for the dynamic map.
3. In the **Priority** field, enter a priority number for the map. Negotiation requests for security associations try to match the highest-priority map first. If that map does not match, the negotiation request continues down the list to the next-highest priority map until a match is made.
4. Click the **Version** drop-down list, and select **v2** to create a map for remote peers using IKEv2.
5. (Optional) Configure Perfect Forward Secrecy (PFS) settings for the dynamic peer by assigning a Diffie-Hellman prime modulus group. PFS provides an additional level of security by ensuring that the IPsec SA key was not derived from any other key, and therefore can not be compromised if another key is broken. Click the **Set PFS** drop-down list and select one of the following groups:
   - Group 1: 768-bit Diffie–Hellman prime modulus group.
   - Group 2: 1024-bit Diffie–Hellman prime modulus group.
   - Group 14: 2048-bit Diffie–Hellman prime modulus group.
   - Group 19: 256-bit random Diffie–Hellman ECP modulus group.
   - Group 20: 384-bit random Diffie–Hellman ECP modulus group.
6. Select the transform set for the map to define a specific encryption and authentication type used by the dynamic peer. Click the **Transform Set** drop-down list, and select the transform set for the dynamic peer.



To view current configuration settings for an IPsec transform-set, access the command-line interface and issue the command **crypto ipsec transform-set tag <transform-set-name>.**

7. Set the **Security Association Lifetime** to define the lifetime of the security association for the dynamic peer, in seconds. The default value is 7200 seconds. To change this value, uncheck the **default** checkbox and enter a value from **300** to **86400** seconds.
8. Click **Done**.

Finalizing WebUI changes

When you have finished configuring your IPsec VPN settings, click **Apply** to apply the new settings before navigating to other pages.

In the CLI

Use the following procedures to use the command-line interface to configure a remote access VPN for L2TP IPsec using IKEv2.

1. Define the server addresses:

```
(host)(config) #vpdn group l2tp
enable
client configuration {dns|wins} <ipaddr1> [<ipaddr2>]
```

2. Enable authentication methods for IKEv2 clients:

```
(host)(config) #crypto isakmp eap-passthrough {eap-mschapv2|eap-peap|eap-tls}
```

3. Create address pools:

```
(host)(config) #ip local pool <pool> <start-ipaddr> <end-ipaddr>
```

4. Configure source NAT:

```
(host)(config) #ip access-list session srcnat user any any src-nat pool <pool> position 1
```

5. If you are configuring a VPN to support machine authentication using certificates, define server certificates for VPN clients using IKEv2:

```
(host)(config) #crypto-local isakmp server-certificate <cert>
```



The IKE pre-shared key value must be between 6-64 characters. To configure a pre-shared IKE key that contains non-alphanumeric characters, surround the key with quotation marks.
For example: **crypto-local isakmp key "key with spaces" fqdn-any**.

6. Define IKEv2 Policies:

```
(host)(config) #crypto isakmp policy <priority>
encryption {3des|aes128|aes192|aes256|des}
version v2
authentication {pre-share|rsa-sig|ecdsa-256ecdsa-384}
group {1|2|19|20}
hash {md5|sha|sha1-96|sha2-256-128|sha2-384-192}
prf PRF-HMAC-MD5|PRF-HMAC-SHA1|PRF-HMAC-SHA256|PRF-HMAC-SHA384
lifetime <seconds>
```

7. Define IPsec Tunnel parameters:

```
(host)(config) #crypto ipsec
mtu <max-mtu>
transform-set <transform-set-name> esp-3des|esp-aes128|esp-aes128-gcm|esp-aes192|esp-
aes256|esp-aes256-gcm|esp-des esp-md5-hmac|esp-null-mac|esp-sha-hmac
```

Configuring a VPN for Smart Card Clients

This section describes how to configure a remote access VPN on the controller for Microsoft L2TP/IPsec clients with smart cards. (A smart card contains a digital certificate which allows user-level authentication without the user entering a username and password.) As described previously in this chapter, L2TP/IPsec requires two levels of authentication: first, IKE SA (machine) authentication, and then user-level authentication with an IKEv2 or PPP-based authentication protocol.

Microsoft clients running Windows 7 (or later versions) support both IKEv1 and IKEv2. Microsoft clients using IKEv2 support machine authentication using RSA certificates (but not ECDSA certificates or pre-shared keys) and smart card user-level authentication with EAP-TLS over IKEv2.



Windows 7 clients without smart cards also support user password authentication using EAP-MSCHAPv2 or PEAP-MSCHAPv2.

Working with Smart Card clients using IKEv2

To configure a VPN for Windows 7 clients using smart cards and IKEv2, follow the procedure described in Configuring a VPN for L2TP/IPsec with IKEv2 in the WebUI on page 346, and ensure that the following settings are configured:

- **L2TP** is enabled
- User Authentication is set to **EAP-TLS**
- IKE version is set to **V2**
- The IKE policy is configured for **ECDSA** or **RSA** certificate authentication

Working with Smart Card Clients using IKEv1

Microsoft clients using IKEv1 (including clients running Windows Vista or earlier versions of Windows) only support machine authentication using a pre-shared key. In this scenario, user-level authentication is performed by an external RADIUS server using PPP EAP-TLS, and client and server certificates are mutually authenticated during the EAP-TLS exchange. During the authentication, the controller encapsulates EAP-TLS messages from the client into RADIUS messages and forwards them to the server.

On the controller, you need to configure the L2TP/IPsec VPN with EAP as the PPP authentication and IKE policy for preshared key authentication of the SA.



On the RADIUS server, you must configure a remote access policy to allow EAP authentication for smart card users and select a server certificate. The user entry in Microsoft Active Directory must be configured for smart cards.

To configure a L2TP/IPsec VPN for clients using smart cards and IKEv1, ensure that the following settings are configured:

1. On a RADIUS server, you must configure a remote access policy to allow EAP authentication for smart card users and select a server certificate. The user entry in Microsoft Active Directory must be configured for smart cards. (For detailed information on creating and managing user roles and policies, see Roles and Policies on page 364.)

- Ensure that RADIUS server is part of the server group used for VPN authentication.
- Configure other VPN settings as described in Configuring a VPN for L2TP/IPsec with IKEv2 in the WebUI on page 346, while selecting the following options:
  - Select **Enable L2TP**
  - Select **EAP** for the Authentication Protocol.
  - Define an IKE Shared Secret to be used for machine authentication. (To make the IKE key global, specify 0.0.0.0 and 0.0.0.0 for both subnet and subnet mask.)
  - Configure the IKE policy for **Pre-Share** authentication.

Configuring a VPN for Clients with User Passwords

This section describes how to configure a remote access VPN on the controller for L2TP/IPsec clients with user passwords. As described previously in this section, L2TP/IPsec requires two levels of authentication: first, IKE

SA authentication, and then user-level authentication with the PAP authentication protocol. IKE SA is authenticated with a preshared key, which you must configure as an IKE shared secret on the controller. User-level authentication is performed by the controller's internal database.

On the controller, you need to configure the following:

- AAA database entries for username and passwords
- VPN authentication profile, which defines the internal server group and the default role assigned to authenticated clients
- L2TP/IPsec VPN with PAP as the PPP authentication (IKEv1 only).
- (For IKEv1 clients) An IKE policy for preshared key authentication of the SA.
- (For IKEv2 clients) A server certificate to authenticate the controller to clients and a CA certificate to authenticate VPN clients.

In the WebUI

Use the following procedure the configure L2TP/IPsec VPN for username/password clients via the WebUI:

1. Navigate to the **Configuration > Security > Authentication > Servers** window.
   a. Select **Internal DB** to display entries for the internal database.
   b. Click **Add User**.
   c. Enter username and password information for the client.
   d. Click **Enabled** to activate this entry on creation.
   e. Click **Apply**.
2. Navigate to the **Configuration > Security > Authentication > L3 Authentication** window.
   a. Under default VPN Authentication Profile, select **Server Group**.
   b. Select the **internal** server group from the drop-down menu.
   c. Click **Apply**.
3. Navigate to the **Configuration > Advanced Services > VPN Services > IPsec** window.
   a. Select **Enable L2TP** (this is enabled by default).
   b. Select **PAP** for Authentication Protocols.
4. Configure other VPN settings as described in Configuring a VPN for L2TP/IPsec with IKEv2 in the WebUI on page 346, while ensuring that the following settings are selected:
   - In the **L2TP and XAUTH Parameters** section of the **Configuration > VPN Services > IPsec** tab, enable **L2TP**.
   - In the **L2TP and XAUTH Parameters** section of the **Configuration > VPN Services > IPsec** tab, select **PAP** as the authentication protocol.

In the CLI

The following example uses the command-line interface to configure a L2TP/IPsec VPN for username/password clients using IKEv1:

```
(host)(config) #vpdn group l2tp
  enable
  ppp authentication pap
  client dns 101.1.1.245

(host)(config) #ip local pool pw-clients 10.1.1.1 10.1.1.250

(host)(config) #crypto isakmp key <key> address 0.0.0.0 netmask 0.0.00

(host)(config) #crypto isakmp policy 1
```

```
authentication pre-share
```

Next, issue the following command in *enable* mode to configure client entries in the internal database:

```
(host)(config) #local-userdb add username <name> password <password>
```

Configuring Remote Access VPNs for XAuth

Extended Authentication (XAuth) is an Internet Draft that allows user authentication after IKE Phase 1 authentication. This authentication prompts the user for a username and password, with user credentials authenticated with an external RADIUS or LDAP server or the controller's internal database. Alternatively, the user can start the client authentication with a smart card, which contains a digital certificate to verify the client credentials. IKE Phase 1 authentication can be done with either an IKE preshared key or digital certificates.

Configuring VPNs for XAuth Clients using Smart Cards

This section describes how to configure a remote access VPN on the controller for Cisco VPN XAuth clients using smart cards. (A smart card contains a digital certificate which allows user-level authentication without the user entering a username and password.) IKE Phase 1 authentication can be done with either an IKE preshared key or digital certificates; for XAuth clients using smart cards, the smart card digital certificates must be used for IKE authentication. The client is authenticated with the internal database on the controller.

On the controller, you need to configure the following:

1. Add entries for Cisco VPN XAuth clients to the controller's internal database, or to an external RADIUS or LDAP server. For details on configuring an authentication server, see Authentication Servers on page 225.



For each client, you need to create an entry in the internal database with the entire Principal name (SubjectAltname in X.509 certificates) or Common Name as it appears on the certificate.

2. Verify that the server with the client data is part of the server group associated with the VPN authentication profile.
3. In the **L2TP and XAUTH Parameters** section of the **Configuration > VPN Services > IPsec** tab, enable **L2TP**.
4. In the **L2TP and XAUTH Parameters** section of the **Configuration > VPN Services > IPsec** tab, enable **XAuth** to enable prompting for the username and password.
5. The Phase 1 IKE exchange for XAuth clients can be either **Main Mode** or **Aggressive Mode**. Aggressive Mode condenses the IKE SA negotiations into three packets (versus six packets for Main Mode). In the **Aggressive Mode** section of the **Configuration > VPN Services > IPsec** tab, Enter the authentication group name for aggressive mode to associate this setting to multiple clients. Make sure that the group name matches the aggressive mode group name configured in the VPN client software.
6. Configure other VPN settings as described in Configuring a VPN for L2TP/IPsec with IKEv2 in the WebUI on page 346, while ensuring that the following settings are selected:
   - In the **L2TP and XAUTH Parameters section** of the **Configuration > VPN Services > IPSEC** tab, enable **L2TP**.
   - n the **L2TP and XAUTH Parameters section** of the **Configuration > VPN Services> IPSEC** tab, enable **XAuth** to enable prompting for the username and password.
   - Define an IKE policy to use **RSA** or **ECDSA** authentication.

The following example describes the steps to use the command-line interface to configure a VPN for Cisco Smart Card Clients using certificate authentication and IKEv1, where the client is authenticated against user entries added to the internal database:

```
(host)(config) #aaa authentication vpn default
```

```
   server-group internal

(host)(config) #no crypto-local isakmp xauth

(host)(config) #vpdn group l2tp
   enable
   client dns 101.1.1.245

(host)(config) #ip local pool sc-clients 10.1.1.1 10.1.1.250

(host)(config) #crypto-local isakmp server-certificate MyServerCert
(host)(config) #crypto-local isakmp ca-certificate TrustedCA

(host)(config) #crypto isakmp policy 1
   authentication rsa-sig
```

Enter the following command in enable mode to configure client entries in the internal database:

```
(host)(config) #local-userdb add username <name> password <password>
```

Configuring a VPN for XAuth Clients Using a Username and Password

This section describes how to configure a remote access VPN on the controller for Cisco VPN XAuth clients using passwords. IKE Phase 1 authentication is done with an IKE preshared key; users are then prompted to enter their username and password, which is verified with the internal database on the controller.

On the controller, you need to configure the following:

1. Add entries for Cisco VPN XAuth clients to the controller's internal database. For details on configuring an authentication server, see Authentication Servers on page 225



For each client, you need to create an entry in the internal database with the entire Principal name (SubjectAltname in X.509 certificates) or Common Name as it appears on the certificate.

2. Verify that the server with the client data is part of the server group associated with the VPN authentication profile.
3. Configure other VPN settings as described in Configuring a VPN for L2TP/IPsec with IKEv2 in the WebUI on page 346, while ensuring that the following settings are selected:
   - In the **L2TP and XAUTH Parameters section** of the **Configuration > VPN Services > IPSEC** tab, enable **L2TP**.
   - In the **L2TP and XAUTH Parameters section** of the **Configuration > VPN Services > IPSEC** tab, enable **XAuth** to enable prompting for the username and password.
   - The IKE policy must have **pre-shared** authentication.

The following example configures a VPN for XAuth IKEv1 clients using a username and passwords. Access the command-line interface and issue the following commands in config mode:

```
(host)(config) #aaa authentication vpn default
   server-group internal

crypto-local isakmp xauth

(host)(config) #vpdn group l2tp
   enable
   client dns 101.1.1.245

(host)(config) #ip local pool pw-clients 10.1.1.1 10.1.1.250

(host)(config) #crypto isakmp key 0987654 address 0.0.0.0 netmask 0.0.00
```

```
(host)(config) #crypto isakmp policy 1
   authentication pre-share
```

Enter the following command in enable mode to configure client entries in the internal database:

```
(host)(config) #local-userdb add username <name> password <password>
```

Working with Remote Access VPNs for PPTP

Point-to-Point Tunneling Protocol (PPTP) is an alternative to L2TP/IPsec. Like L2TP/IPsec, PPTP provides a logical transport mechanism to send PPP frames and tunneling or encapsulation, so that the PPP frames can be sent across an IP network. PPTP relies on the PPP connection process to perform user authentication and protocol configuration.

With PPTP, data encryption begins after PPP authentication and connection process is completed. PPTP connections use Microsoft Point-to-Point Encryption (MPPE), which uses the Rivest-Shamir-Aldeman (RSA) RC-4 encryption algorithm. PPTP connections require user-level authentication through a PPP-based authentication protocol (MSCHAPv2 is the currently-supported method).

In the WebUI

1. Navigate to the **Configuration > Advanced Services > VPN Services > PPTP**page.
2. To enable PPTP, select **Enable PPTP**.
3. Select either **MSCHAP** or **MSCHAPv2** as the authentication protocol.
4. Configure IP addresses of the primary and secondary DNS servers.
5. Configure the primary and secondary WINS Server IP addresses that are pushed to the VPN Dialer.
6. Configure the VPN Address Pool.
   a. Click **Add**. The **Add Address Pool** window displays.
   b. Specify the pool name, start address, and end address.
   c. Click **Done**.
7. Click **Apply** to apply the changes made before navigating to other pages.

In the CLI

```
(host)(config) #vpdn group pptp
   enable
   client configuration {dns|wins} <ipaddr1> [<ipaddr2>]
   ppp authentication {mschapv2}
(host)(config) #pptp ip local pool <pool> <start-ipaddr> <end-ipaddr>
```

Working with Site-to-Site VPNs

Site-to-site VPN allows sites at different physical locations to securely communicate with each other over a Layer-3 network such as the Internet. You can use Dell controllers instead of VPN concentrators to connect the sites. You can also use a VPN concentrator at one site and a controller at the other site.

The Dell controller supports the following IKE SA authentication methods for site-to-site VPNs:

- Preshared key: Note that the same IKE shared secret must be configured on both the local and remote sites.
- Suite-B cryptographic algorithms
- Digital certificates: You can configure a RSA or ECDSA server certificate and a CA certificate for each site-to-site VPN IPsec map configuration. If you use certificate-based authentication, the peer must be identified by its certificate subject-name distinguished name (for deployments using IKEv2) or by the peer's IP address

(for IKEv1). For more information about importing server and CA certificates into the controller, see Management Access on page 778.

NOTE

Certificate-based authentication is only supported for site-to-site VPN between two controllers with static IP addresses. IKEv1 site-to-site tunnels can not be created between master and local controllers.

Working with Third-Party Devices

Dell controllers can use IKEv1 or IKEv2 to establish a site-to-site VPN between another Dell controller or between that controller and third-party remote client devices. Devices running Microsoft® Windows 2008 can use Suite-B cryptographic algorithms and IKEv1 to support authentication using RSA or ECDSA. strongSwan® 4.3 devices can use IKEv2 to support authentication using RSA or ECDSA certificates, Suite-B cryptographic algorithms, and pre-shared keys. These two remote clients are tested to work with Dell controllers using Suite-B cryptographic algorithm.

Working with Site-to-Site VPNs with Dynamic IP Addresses

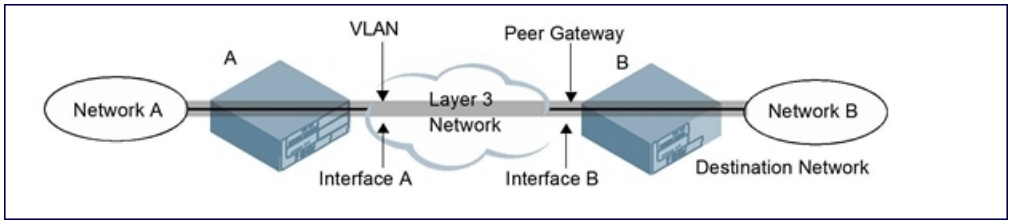
ArubaOS supports site-to-site VPNs with two statically addressed controllers, or with one static and one dynamically addressed controller. Two methods are supported to enable dynamically addressed peers:

- Pre-shared Key Authentication with IKE Aggressive Mode: The Dell controller with a dynamic IP address must be configured to be the initiator of IKE Aggressive-mode for Site-Site VPN, while the controller with a static IP address must be configured as the responder of IKE Aggressive mode. Note that when the controller is operating in FIPS mode, IKE aggressive mode must be disabled.
- X.509 certificates: IPsec peers will identify each other using the subject name of X.509 certificates. IKE operates in main mode when this option is selected. This method is preferred from a security perspective.

Understanding VPN Topologies

You must configure VPN settings on the controllers at both the local and remote sites. In the following figure, a VPN tunnel connects Network A to Network B across the Internet.

Figure 36 *Site-to-Site VPN Configuration Components*



To configure the VPN tunnel on controller A, you need to configure the following:

- The source network (Network A)
- The destination network (Network B)
- The VLAN on which controller A's interface to the Layer-3 network is located (Interface A in Figure 36)
- The peer gateway, which is the IP address of controller B's interface to the Layer-3 network (Interface B in Figure 36)

NOTE

Configure VPN settings on the controllers at both the local and remote sites.

Configuring Site-to-Site VPNs

Use the following procedures to create a site-to-site VPN via the WebUI or command-line interfaces.

In the WebUI

1. Navigate to the **Configuration > Advanced Services > VPN Services > Site-to-Site** page.
2. In the **IPsec Maps** section, click **Add** to open the **Add IPsec Map** window.
3. Enter a name for this VPN connection in the **Name** field.
4. Enter a priority level for the IPsec map. Negotiation requests for security associations try to match the highest-priority map first. If that map does not match, the negotiation request continues down the list to the next-highest priority map until a match is made.
5. In the **Source Network** and **Source Subnet Mask fields**, enter the IP address and netmask for the source (the local network connected to the controller). (See controller A in Figure 36.)
6. In the **Destination Network** and **Destination Subnet Mask** fields, enter the IP address and netmask for the destination (the remote network to which the local network communicates). (See controller B in Figure 36.)
7. If you use IKEv1 to establish a site-to-site VPN to a statically addressed remote peer, in the **Peer Gateway** field, enter the IP address of the interface used by remote peer to connect to the L3 network. (See Interface B in Figure 36.) If you are configuring an IPsec map for a dynamically addressed remote peer, you must leave the peer gateway set to its default value of **0.0.0.0**.
8. If you use IKEv2 to establish a site-to-site VPN to a statically addressed remote peer, identify the peer device by entering its certificate subject name in the **Peer Certificate Subject Name** field.



To identify the subject name of a peer certificate, access the command-line interface and issue the command `show crypto-local pki servercert <certname> subject`

9. The **Security Association Lifetime** parameter defines the lifetime of the security association, in seconds and kilobytes. For seconds, the default value is **7200**. To change this value, uncheck the **default** checkbox and enter a value from 300 to 86400 seconds. Range: 1000–1000000000 kilobytes.
10. Click the **Version** drop-down list and select **V1** to configure the VPN for IKEv1, or **V2** for IKEv2.
11. Select the **VLAN** that contains the interface of the local controller which connects to the Layer-3 network. (See Interface A in Figure 36.)

    This determines the source IP address used to initiate IKE. If you select **0** or **None**, the default is the VLAN of the controller's IP address (either the VLAN where the loopback IP is configured, or VLAN 1 if no loopback IP is configured).
12. If you enable **Perfect Forward Secrecy (PFS)** mode, new session keys are not derived from previously used session keys. Therefore, if a key is compromised, that compromised key does not affect any previous session keys. PFS mode is disabled by default. To enable this feature, click the **PFS** drop-down list and select one of the following **Perfect Forward Secrecy** modes:
    - **group1**: 768-bit Diffie–Hellman prime modulus group.
    - **group2**: 1024-bit Diffie–Hellman prime modulus group.
    - **group 14**: 2048-bit Diffie–Hellman prime modulus group.
    - **group19**: 256-bit random Diffie–Hellman ECP modulus group.
    - **group20**: 384-bit random Diffie–Hellman ECP modulus group.
13. Select **Pre-Connect** to have the VPN connection established, even if there is no traffic being sent from the local network. If you do not select this the VPN connection established only when traffic is sent from the local network to the remote network.

14. Select **Trusted Tunnel** if traffic between the networks is trusted. If you do not select this, traffic between the networks is untrusted.
15. Select the **Enforce NATT** checkbox to always enforce UDP 4500 for IKE and IPSEC. This option is disabled by default.
16. Add one or more transform sets to be used by the IPsec map. Click the **Transform Set** drop down list, select an existing transform set, then click the arrow button by the drop-down list to add that transform set to the IPsec map.
17. For site-to-site VPNs with dynamically addressed peers, click the **Dynamically Addressed Peers** checkbox.
    a. Select **Initiator** if the dynamically addressed switch is the *initiator* of IKE Aggressive-mode for Site-Site VPN, or select **Responder** if the dynamically addressed switch is the *responder* for IKE Aggressive-mode.
    b. In the **FQDN** field, enter a fully qualified domain name (FQDN) for the controller. If the controller is defined as a dynamically addressed responder, you can select **all peers** to make the controller a responder for all VPN peers, or select **Per Peer ID** and specify the FQDN to make the controller a responder for one specific initiator only.
18. Select one of the following authentication types:
    a. For pre-shared key authentication, select **Pre-Shared Key**, then enter a shared secret in the **IKE Shared Secret** and **Verify IKE Shared Secret** fields. This authentication type is generally required in IPsec maps for a VPN with a dynamically addressed peers but can also be used for a static site-to-site VPN.
    b. For certificate authentication, select **Certificate**, then click the **Server Certificate** and **CA certificate** drop-down lists to select certificates previously imported into the controller. See Management Access on page 778 for more information.
19. Click **Done** to apply the site-to-site VPN configuration.
20. Click **Apply**.
21. Click the **IPSEC** tab to configure an IKE policy.
    a. Under IKE Policies, click **Add**to open the **IPSEC Add Policy** configuration page.
    b. Set the **Priority** to **1** for this configuration to take priority over the Default setting.
    c. Set the **Version type** to match the IKE version you selected in Step 10 above.
    d. Set the **Encryption type** from the drop-down list**.**
    e. Set the **HASH Algorithm** from the drop-down list.
    f. Set the Authentication to **PRE-SHARE** if you use preshared keys. If you use certificate-based IKE, select **RSA** or **ECDSA**.
    g. Set the **Diffie–Hellman Group** from the drop-down list.
    h. The IKE policy selections, including any preshared key, need to be reflected in the VPN client configuration. When using a third party VPN client, set the VPN configuration on clients to match the choices made above. If you use the Dell dialer, you must configure the dialer prior to downloading the dialer onto the local client.
    i. Click **Done** to activate the changes.
    j. Click **Apply.**

In the CLI

To use the command-line interface to configure a site-to-site VPN with two static IP controllers using IKEv1, issue the following commands:

```
(host)(config) #crypto-local ipsec-map <name> <priority>
   src-net <ipaddr> <mask>
   dst-net <ipaddr> <mask>
   peer-ip <ipaddr>
   vlan <id>
```

```
   version v1|v2
   peer-cert-dn <peer-dn>
   pre-connect enable|disable
   trusted enable
```

For certificate authentication:

```
   set ca-certificate <cacert-name>
   set server-certificate <cert-name>

(host)(config) #crypto isakmp policy <priority>
   encryption {3des|aes128|aes192|aes256|des}
   version v1|v2
   authentication {rsa-sig|ecdsa-256ecdsa-384}
   group {1|2|19|20}
   hash {md5|sha|sha1-96|sha2-256-128|sha2-384-192}
   lifetime <seconds>
```

For preshared key authentication:

```
(host)(config) #crypto-local isakmp key <key> address <ipaddr> netmask <mask>

(host)(config) #crypto isakmp policy <priority>
   encryption {3des|aes128|aes192|aes256|des}
   version v1|v2
   authentication pre-share
   group {1|2|19|20}
   hash {md5|sha|sha1-96|sha2-256-128|sha2-384-192}
   lifetime <seconds>
```

To configure site-to-site VPN with a static and a dynamically addressed controller that initiates IKE Aggressive-mode for Site-Site VPN:

```
(host)(config) #crypto-local ipsec-map <name> <priority>
   src-net <ipaddr> <mask>
   dst-net <ipaddr> <mask>
   peer-ip <ipaddr>local-fqdn <local_id_fqdn>
   vlan <id>
   pre-connect enable|disable
   trusted enable
```

For the Pre-shared-key:

```
(host)(config) #crypto-local isakmp key <key> address <ipaddr> netmask 255.255.255.255
```

For a static IP controller that responds to IKE Aggressive-mode for Site-Site VPN:

```
crypto-local ipsec-map <name2> <priority>
   src-net <ipaddr> <mask>
   dst-net <ipaddr> <mask>
   peer-ip 0.0.0.0
   peer-fqdn fqdn-id <peer_id_fqdn>
   vlan <id>
   trusted enable
```

For the Pre-shared-key:

```
(host)(config) #crypto-local isakmp key <key> fqdn <fqdn-id>
```

For a static IP controller that responds to IKE Aggressive-mode for Site-Site VPN with One PSK for All FQDNs:

```
(host)(config) #crypto-local ipsec-map <name2> <priority>
   src-net <ipaddr> <mask>
```

```
peer-ip 0.0.0.0
peer-fqdn any-fqdn
vlan <id>
trusted enable
```

For the Pre-shared-key for All FQDNs:

```
(host)(config) #crypto-local isakmp key <key> fqdn-any
```

Detecting Dead Peers

Dead Peer Detection (DPD) is enabled by default on the controller for site-to-site VPNs. DPD, as described in RFC 3706, "A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers," uses IPsec traffic patterns to minimize the number of IKE messages required to determine the liveliness of an IKE peer.

After a dead peer is detected, the controller tears down the IPsec session. Once the network path or other failure condition has been corrected, a new IPsec session is automatically re-established.

To configure DPD parameters, issue the following commands via the command-line interface:

```
(host)(config) #crypto-local isakmp dpd idle-timeout <idle_seconds> retry-timeout <retry_
seconds> retry-attempts <number>
```

About Default IKE Policies

ArubaOS includes the following default IKE policies. These policies are predefined and canbe edited and deleted. You can do this by using the **crypto isakmp policy** and **crypto dynamic-map** CLI commands. Or, use the WebUI and navigate to **Advanced Services > VPN Services > IPSEC** and use the **Delete** button next to the default IKE policy or IPsec dynamic map you want to delete.

Table 60: *Default IKE Policy Settings*

Policy Name	Policy Number	IKE Version	Encryption Algorithm	Hash Algorithm	Authentica-tion Method	PRF Method	Diffie-Hellman Group
Default protection suite	10001	IKEv1	3DES-168	SHA 160	Pre-Shared Key	N/A	2 (1024 bit)
Default RAP Certificate protection suite	10002	IKEv1	AES -256	SHA 160	RSA Signature	N/A	2 (1024 bit)
Default RAP PSK protection suite	10003		AES -256	SHA 160	Pre-Shared Key	N/A	2 (1024 bit)
Default RAP IKEv2 RSA protection suite	1004	IKEv2	AES -256	SSHA160	RSA Signature	hmac-sha1	2 (1024 bit)

Policy Name	Policy Number	IKE Version	Encryption Algorithm	Hash Algorithm	Authentica-tion Method	PRF Method	Diffie-Hellman Group
Default Cluster PSK protection suite	10005	IKEv1	AES -256	SHA160	Pre-Shared Key	Pre-Shared Key	2 (1024 bit)
Default IKEv2 RSA protection suite	1006	IKEv2	AES - 128	SHA 96	RSA Signature	hmac-sha1	2 (1024 bit)
Default IKEv2 PSK protection suite	10007	IKEv2	AES - 128	SHA 96	Pre-shared key	hmac-sha1	2 (1024 bit)
Default Suite-B 128bit ECDSA protection suite	10008	IKEv2	AES - 128	SHA 256-128	ECDSA-256 Signature	hmac-sha2-256	Random ECP Group (256 bit)
Default Suite-B 256 bit ECDSA protection suite	10009	IKEv2	AES -256	SHA 384-192	ECDSA-384 Signature	hmac-sha2-384	Random ECP Group (384 bit)
Default Suite-B 128bit IKEv1 ECDSA protection suite	10010	IKEv1	AES-GCM-128	SHA 256-128	ECDSA-256 Signature	hmac-sha2-256	Random ECP Group (256 bit)
Default Suite-B 256-bit IKEv1 ECDSA protection suite	10011	IKEv1	AES-GCM-256	SHA 256-128	ECDSA-256 Signature	hmac-sha2-256	Random ECP Group (256 bit)

Working with VPN Dialer

For Windows clients, a dialer can be downloaded from the controller to auto configure tunnel settings on the client.

Configuring VPN Dialer

Use the following procedures to configure the VPN dialer via the WebUI or command-line interfaces

In the WebUI

1. Navigate to the **Configuration > Advanced Services > VPN Services > Dialers** page. Click **Add** to add a new dialer or the **Edit** tab to edit an existing dialer.
2. Enter the Dialer Name that identifies this setting.
3. Configure the dialer to work with PPTP or L2TP by selecting **Enable PPTP** or **Enable L2TP**.
4. Select the authentication protocol. This should match the L2TP or PPTP authentication type configured for the VPN in the **Configuration > Advanced Services > VPN Services > IPSEC** window.
5. (Optional) Select **Send Direct Network Traffic In Clear** to enable "split tunneling" functionality so that traffic destined for the internal network is tunneled, while traffic for the Internet is not.
   This option is not recommended for security reasons.
6. (Optional) Select **Disable Wireless Devices When Client is Wired** to allow the dialer to shut down the wireless interface when it detects that a wired network connection is in use.
7. (Optional) Select **Enable SecurID New and Next Pin Mode** to enable site-to-site VPN support for SecurID new and next pin modes.
8. For L2TP:
   - Set the IKE Hash Algorithm to the value defined in the IKE policy on the **Advanced Services > VPN Services > IPSEC** window.
   - If a preshared key is configured for an IKE Shared Secret in the **VPN Services > IPSEC** window, enter the key.
   - The key you enter in the **Dialers** window must match the preshared key configured on the IPsec page.
   - Select the IPsec Mode Group that matches the Diffie–Hellman Group configured for the IPsec policy.
   - Select the IPsec Encryption that matches the encryption configured for the IPsec policy.
   - Select the IPsec Hash Algorithm that matches the hash algorithm configured for the IPsec policy.
9. Click **Done** to apply the changes made prior to navigating to another page.

In the CLI

Issue the following commands to configure the VPN dialer via the CLI:

```
(host(config) #vpn-dialer <name>
   enable {dnctclear|l2tp|pptp|secureid_newpinmode|wirednowifi}
   ike authentication {pre-share <key>|rsa-sig}
   ike encryption {3des|des}
   ike group {1|2}
   ike hash {md5|sha}
   ipsec encryption {esp-3des|esp-des}
   ipsec hash {esp-md5-hmac|esp-sha-hmac}
   ppp authentication {cache-securid|chap|mschap|mschapv2|pap}
```

Assigning a Dialer to a User Role

The VPN dialer can be downloaded using Captive Portal. For the user role assigned through Captive Portal, configure the dialer by the name used to identify the dialer.

For example, if the captive portal client is assigned the *guest* role after logging on through captive portal and the dialer is called *mydialer*, configure *mydialer* as the dialer to be used in the guest role.

In the WebUI

1. Navigate to the **Configuration > Security > Access Control > User Roles** page.
2. Click **Edit** for the user role.
3. Under VPN Dialer, select the dialer you configured and click **Change**.
4. Click **Apply**.

In the CLI

To configure the captive portal dialer for a user role via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host) (config) #user-role <role>
   dialer <name>
```

The client in a Dell user-centric network is associated with a *user role*, which determines the client's network privileges, how often it must re-authenticate, and which bandwidth contracts are applicable. A *policy* is a set of rules that applies to traffic that passes through the Dell controller. You specify one or more policies for a user role. Finally, you can assign a user role to clients before or after they authenticate to the system.

This chapter describes assigning and creating roles and policies using the ArubaOS CLI or WebUI. Roles and policies can also be configured for WLANs associated with the "default" ap-group via the WLAN Wizard: **Configuration > Wizards > WLAN Wizard**. Follow the steps in the workflow pane within the wizard and refer to the help tab for assistance.

Topics in this chapter include:

- Configuring Firewall Policies on page 364
- Creating a Firewall Policy on page 365
- Creating a Network Service Alias on page 368
- Creating an ACL White List on page 369
- User Roles on page 370
- Assigning User Roles on page 372
- Understanding Global Firewall Parameters on page 377
- Using AppRF 2.0 on page 381



This chapter describes configuring firewall policies and parameters that relate to IPv4 traffic. See IPv6 Support on page 175 for information about configuring IPv6 firewall policies and parameters.

Configuring Firewall Policies

A firewall policy identifies specific characteristics about a data packet passing through the Dell controller and takes some action based on that identification. In a Dell controller, that action can be a firewall-type action such as permitting or denying the packet, an administrative action such as logging the packet, or a quality of service (QoS) action such as setting 802.1p bits or placing the packet into a priority queue. You can apply firewall policies to user roles to give differential treatment to different users on the same network, or to physical ports to apply the same policy to all traffic through the port.

Firewall policies differ from access control lists (ACLs) in the following ways:

- Firewall policies are *stateful*, meaning that they recognize flows in a network and keep track of the state of sessions. For example, if a firewall policy permits telnet traffic from a client, the policy also recognizes that inbound traffic associated with that session should be allowed.
- Firewall policies are *bi-directional*, meaning that they keep track of data connections traveling into or out of the network. ACLs are normally applied to either traffic inbound to an interface or outbound from an interface.
- Firewall policies are *dynamic*, meaning that address information in the policy rules can change as the policies are applied to users. For example, the alias *user* in a policy automatically applies to the IP address assigned to a particular user. ACLs typically require static IP addresses in the rule.



You can apply IPv4 and IPv6 firewall policies to the same user role. See IPv6 Support on page 175 for information about configuring IPv6 firewall policies.

Working With Access Control Lists (ACLs)

Access control lists (ACLs) are a common way of restricting certain types of traffic on a physical port. ArubaOS provides the following types of ACLs:

- Standard ACLs permit or deny traffic based on the source IP address of the packet. Standard ACLS can be either named or numbered, with valid numbers in the range of 1-99 and 1300-1399. Standard ACLs use a bitwise mask to specify the portion of the source IP address to be matched.
- Extended ACLs permit or deny traffic based on source or destination IP address, source or destination port number, or IP protocol. Extended ACLs can be named or numbered, with valid numbers in the range 100-199 and 2000-2699.
- MAC ACLs are used to filter traffic on a specific source MAC address or range of MAC addresses. Optionally, you can mirror packets to a datapath or remote destination for troubleshooting and debugging purposes. MAC ACLs can be either named or numbered, with valid numbers in the range of 700-799 and 1200-1299.
- Ethertype ACLs are used to filter based on the Ethertype field in the frame header. Optionally, you can mirror packets to a datapath or remote destination for troubleshooting and debugging purposes. Ethertype ACLs can be either named or numbered, with valid numbers in the range of 200-299.These ACLs can be used to permit IP while blocking other non-IP protocols, such as IPX or AppleTalk.
- Service ACLs provide a generic way to restrict how protocols and services from specific hosts and subnets to the controller are used. Rules with this ACL are applied to all traffic on the controller regardless of the ingress port or VLAN.

ArubaOS provides both standard and extended ACLs for compatibility with router software from popular vendors, however firewall policies provide equivalent and greater function than standard and extended ACLs and should be used instead.

You can apply MAC and Ethertype ACLs to a user role, however these ACLs only apply to non-IP traffic *from* the user.

Support for Desktop Virtualization Protocols

ArubaOS supports desktop virtualization protocols by providing preconfigured ACLs for Citrix and VMware clients. You can apply these ACLs to the user-role when using the Virtual Desktop Infrastructure (VDI) clients. This ensures that any enterprise application that uses the VDI client performs optimally with appropriate QoS.



Disable the voice aware ARM when applying the ACLs for the VDI clients as the virtual desktop sessions may prevent the ARM scanning.

Creating a Firewall Policy

This section describes how to configure the rules that constitute a firewall policy. A firewall policy can then be applied to a user role (until the policy is applied to a user role, it does not have any effect).

Table 61 describes required and optional parameters for a rule.

Table 61: *Firewall Policy Rule Parameters*

Field	Description
Source (required)	Source of the traffic, which can be one of the following: • any: Acts as a wildcard and applies to any source address. • user: This refers to traffic from the wireless client. • host: This refers to traffic from a specific host. When this option is chosen, you must configure the IP address of the host. • network: This refers to a traffic that has a source IP from a subnet of IP addresses. When this option is chosen, you must configure the IP address and network mask of the subnet. • alias: This refers to using an alias for a host or network. You configure the alias by navigating to the **Configuration > Advanced Services > Stateful Firewall > Destination** page.
Destination (required)	Destination of the traffic, which can be configured in the same manner as Source.
Service (required)	Type of traffic, which can be one of the following: • **any**: This option specifies that this rule applies to any type of traffic. • tcp: Using this option, you configure a range of TCP port(s) to match for the rule to be applied. • udp: Using this option, you configure a range of UDP port(s) to match for the rule to be applied. • service: Using this option, you use one of the pre-defined services (common protocols such as HTTPS, HTTP, and others) as the protocol to match for the rule to be applied. You can also specify a network service that you configure by navigating to the **Configuration > Advanced Services > Stateful Firewall > Network Services** page. • protocol: Using this option, you specify a different layer 4 protocol (other than TCP/UDP) by configuring the IP protocol value.
Action (required)	The action that you want the controller to perform on a packet that matches the specified criteria. This can be one of the following: • permit: Permits traffic matching this rule. • drop: Drops packets matching this rule without any notification. • reject: Drops the packet and sends an ICMP notification to the traffic source. • src-nat: Performs network address translation (NAT) on packets matching the rule. When this option is selected, you need to select a NAT pool. (If this pool is not configured, you configure a NAT pool by navigating to the **Configuration > Advanced > Security > Advanced > NAT Pools**). Source IP changes to the outgoing interface IP address (implied NAT pool) or from the pool configured (manual NAT pool). This action functions in tunnel/decrypt-tunnel forwarding mode. • dst-nat: This option redirects traffic to the configured IP address and destination port. An example of this option is to redirect all HTTP packets to the captive portal port on the Dell controller as used in the pre-defined policy called “captiveportal”. This action functions in tunnel/decrypt-tunnel forwarding mode. User should configure the NAT pool in the controller. • dual-nat: This option performs both source and destination NAT on packets matching the rule. Forward packets from source network to destination; re-mark them with destination IP of the target network. This action functions in tunnel/decrypt-tunnel forwarding mode. User should configure the NAT pool in the controller. • redirect to tunnel: This option redirects traffic into a GRE tunnel. This option is used primarily to redirect all guest traffic into a GRE tunnel to a DMZ router/switch. • redirect to ESI group: This option redirects traffic to the specified ESI server group. You also specify the direction of traffic to be redirected: forward, reverse, or both directions. • route**:** Specify the next hop to which packets are routed, which can be one of the following:

Field	Description
	• dst-nat: Destination IP changes to the IP configured from the NAT pool. This action functions in bridge/split-tunnel forwarding mode. User should configure the NAT pool in the controller. • src-nat:Source IP changes to RAP's external IP. This action functions in bridge/split-tunnel forwarding mode and uses implied NAT pool.
Log (optional)	Logs a match to this rule. This is recommended when a rule indicates a security breach, such as a data packet on a policy that is meant only to be used for voice calls.
Mirror (optional)	Mirrors session packets to datapath or remote destination.
Queue (optional)	The queue in which a packet matching this rule should be placed. Select **High** for higher priority data, such as voice, and **Low** for lower priority traffic.
Time Range (optional)	Time range for which this rule is applicable. Configure time ranges on the **Configuration > Security > Access Control > Time Ranges** page.
Pause ARM Scanning (optional)	Pause ARM scanning while traffic is present. Note that you must enable "VoIP Aware Scanning" in the ARM profile for this feature to work.
Black List (optional)	Automatically blacklists a client that is the source or destination of traffic matching this rule. This option is recommended for rules that indicate a security breach where the blacklisting option can be used to prevent access to clients that are attempting to breach the security.
White List (optional)	A rule must explicitly permit a traffic session before it is forwarded to the controller. The last rule in the white list denies everything else. Configure white list ACLs on the **Configuration > Advanced Services> Stateful Firewall> White List (ACL)** page.
TOS (optional)	Value of type of service (TOS) bits to be marked in the IP header of a packet matching this rule when it leaves the controller.
802.1p Priority (optional)	Value of 802.1p priority bits to be marked in the frame of a packet matching this rule when it leaves the controller.

The following example creates a policy 'web-only' that allows web (HTTP and HTTPS) access.

In the WebUI

1. Navigate to the **Configuration > Security > Access Control > Policies** page on the WebUI.
2. To configure a firewall policy, select the policy type from the Policies title bar. You can select **All**, **IPv4 Session**, **IPv6 Session**, **Ethernet**, **MAC**, **Standard** or **Extended**.
3. Click **Add** to create a new policy.
4. If you selected All in Step 2, then select the type of policy you are adding from the **Policy Type** drop-down menu.
5. Click **Add** to add a rule that allows HTTP traffic.
   a. Under Service, select service from the drop-down list.
   b. Select svc-http from the scrolling list.
   c. Click **Add**.
6. Click **Add** to add a rule that allows HTTPS traffic.

a. Under Service, select service from the drop-down list.
b. Select svc-https from the scrolling list.
c. Click **Add**.



Rules can be re-ordered by using the up and down buttons provided for each rule.

7. Click **Apply** to apply this configuration. The policy is not created until the configuration is applied.

In the CLI

```
(host)(config) #ip access-list session web-only
   any any svc-http permit
   any any svc-https permit
```

Creating a Network Service Alias

A network service alias defines a TCP, UDP or IP protocol and a list or range of ports supported by that service. When you create a network service alias, you can use that alias when specifying the network service for multiple session ACLs.

In the WebUI

1. Navigate to the **Configuration > Advanced Services> Stateful Firewall > Network Services** page on the WebUI.
2. Click **Add** to create a new alias.
3. Enter a name for the alias in the **Service Name** field.
4. In the **Protocol** section, select either TCP or UDP, or select Protocol and enter the IP protocol number of the protocol for which you want to create an alias.
5. In the **Port Type** section, specify whether you want to define the port by a contiguous range of ports, or by a list of non-contiguous port numbers.
   - If you selected **Range**, enter the starting and ending port numbers in the **Starting Port** and **End Port** fields.
   - If you selected list, enter a comma-separated list of port numbers.
6. To limit the service alias to a specific application, click the **Application Level Gateway (ALG)** drop-down list and select one of the following service types
   - dhcp: Service is DHCP
   - dns: Service is DNS
   - ftp: Service is FTP
   - h323: Service is H323
   - noe: Service is Alcatel NOE
   - rtsp: Service is RTSP
   - sccp: Service is SCCP
   - sip: Service is SIP
   - sips: Service is Secure SIP
   - svp: Service is SVP
   - tftp: Service is TFTP
   - vocera: Service is VOCERA
7. Click **Apply** to save your changes.

In the CLI

To define a service alias via the command-line interface, access the CLI in config mode and issue the following command:

```
(host)(config) #netservice <name> <protocol>|tcp|udp {list <port>,<port>}|{<port> [<port>]}
[ALG <service>]
```

Creating an ACL White List

The ACL White List consists of rules that explicitly permit or deny session traffic from being forwarded to or blocked from the controller. The white list protects the controller during traffic session processing by prohibiting traffic from being automatically forwarded to the controller if it was not specifically denied in a blacklist. The maximum number of entries allowed in the ACL White List is 64. To create an ACL white list, you must first define a white list bandwidth contract, and then assign it to an ACL.

In the WebUI

1. Navigate to the **Configuration > Advanced Services > Stateful Firewall > White List BW Contracts** page.
2. Click **Add** to create a new contract.
3. In the **White list contract name** field, enter the name of a bandwidth contract.
4. The **Bandwidth Rate** field allows you to define a bandwidth rate in either kbps or Mbps. Enter a rate value the **Bandwidth rate** field, then click the drop-down list and select either kbps or Mbps.
5. Click **Done**.

Configuring the ACL White List in the WebUI

1. Navigate to the **Configuration > Stateful Firewall> ACL White List**page.
2. To add an entry, click the **Add**button at the bottom of the page. The **Add New Protocol**section displays.
3. Click the **Action** drop-down list and select **Permit or Deny. Permit** allows session traffic to be forwarded to the controller while **Deny** blocks session traffic.
4. Click the IP Version drop-down list and select the**IPv4** or IPv6 filter. You need to select one of three following choices from the **Source** drop-down list:
   - For a specific IPv4 or IPv6 filter, select **IP/Mask**. Enter the IP address and mask of the IPv4 or IPv6 filter in the corresponding fields.
   - For a IPv4 or IPv6 host, select **Any** and enter the source address.
5. In the **IP Protocol Number** or **IP Protocol** field, enter the number for a protocol or select the protocol from the drop-down list used by session traffic.
6. In the **Starting Port**s field, enter a starting port. This is the first port, in the port range, on which permitted or denied session traffic is running. Port range: 1–65535.
7. In the **End Ports** field, enter an ending port. This is the last port, in the port range, on which permitted or denied session traffic is running. Port range: 1–65535.
8. (Optional) Click the **White list Bandwidth Contract** drop-down list and specify the name of a bandwidth contract to apply to the session traffic. For further information on creating Bandwidth Contracts, see User Roles on page 370
9. Click **Done**. The ACL displays on the white list section.
10. To delete an entry, click **Delete** next to the entry you want to delete.
11. Click **Apply** to save changes.

Configuring the White List Bandwidth Contract in the CLI

```
(host)(config) #cp-bandwidth-contract <name> {mbits <1..2000>}|{kbits <256..2000000>}
```

Configuring the ACL White List in the CLI

Use the following CLI command to create ACL White Lists.

```
(host) (config-fw-cp)ipv4|ipv6 deny|permit <ip-addr><ip-mask>|any|{host <ip-addr>} proto{<ip-protocol-number> ports <start port number> <end port number>}
|ftp|http|https|icmp|snmp|ssh|telnet|tftp [bandwidth-contract <name>]
```

To create a whitelist ACL that allows traffic on an ipv4 filter with the ipv4 source address 10.10.10.10 and the ipv4 source mask 2.2.2.2 where the protocol is ftp and the the bandwidth contract name is mycontract.

```
(host) (config-fw-cp) #ipv4 permit 10.10.10.10 2.2.2.2 proto ftp bandwidth-contract name mycontract
```

to create a whitelist ACL entry that denies traffic using protocol 2 on port 5000 from being forwarded to the controller:

```
(host) (config-fw-cp) deny proto 2 ports 5000 5000
```

User Roles

User roles are comprised of user role settings, firewall policies, and bandwidth contracts. This section describes the procedure to create a new user role, and associate a firewall policy with that role.

This section describes how to create a new user role. When you create a user role, you must specify one or more firewall policies for the role.

In the WebUI

1. Navigate to the **Configuration > Security > Access Control > User Roles** page.
2. Click **Add** to create and configure a new user role.
3. Enter a user role name.
4. Under Firewall Policies, click **Add**.
5. Select one of the following three options to add a policy to the role.
   - To use an associate an existing policy to the user role, select **Choose from Configured Policies** then select an existing policy from the drop-down list.
   - to create a new policy based upon the settings of an existing policy, select **Create New Policy from Existing Policy** drop-down list, then select an existing policy from the drop-down list. The **Policies** page appears, allowing you to configure a new firewall policy.
   - To create and configure an entirely new policy, select **Create New Policy**, then click **Create**. The **Policies** page appears, allowing you to configure a new firewall policy.
6. Click **Done** to add the policy to the user role.
7. (Optional) If the user role contains more than one firewall policy, use the up and down arrows to assign priorities to each role. The higher the policy on the list, the higher its priority.
8. In the **Misc. Configuration** section, enter configuration values as described in Table 62.
9. Click **Apply.**
10. Next, you must assign the user role to a AAA profile. After assigning the user role you can use the **show reference user-role <role>** command to see the profiles that reference this user role.For more information, see Assigning User Roles on page 372

Table 62: *User Role Parameters*

Field	Description
Role name	Name of the user role
Re-authentication Interval (optional)	Time, in minutes, after which the client is required to reauthenticate. Enter a value between 0-4096. 0 disables reauthentication. Default: 0 (disabled)
Role VLAN ID (optional)	By default, a client is assigned a VLAN on the basis of the ingress VLAN for the client to the controller. You can override this assignment and configure the VLAN ID that is to be assigned to the user role. You configure a VLAN by navigating to the **Configuration > Network > VLANs** page.
Bandwidth Contract (optional)	You can assign a bandwidth contract to provide an upper limit to upstream or downstream bandwidth utilized by clients in this role. You can select the Per User option to apply the bandwidth contracts on a per-user basis instead of to all clients in the role. For more information, see User Roles on page 370.
VPN Dialer (optional)	This assigns a VPN dialer to a user role. For details about VPN dialer, see Virtual Private Networks on page 337. Select a dialer from the drop-down list and assign it to the user role. This dialer will be available for download when a client logs in using captive portal and is assigned this role.
L2TP Pool (optional)	This assigns an L2TP pool to the user role. For more details about L2TP pools, see Virtual Private Networks on page 337. Select the required L2TP pool from the list to assign to the user role. The inner IP addresses of VPN tunnels using L2TP will be assigned from this pool of IP addresses for clients in this user role.
PPTP Pool (optional)	This assigns a PPTP pool to the user role. For more details about PPTP pools, see Virtual Private Networks on page 337. Select the required PPTP pool from the list to assign to the user role. The inner IP addresses of VPN tunnels using PPTP will be assigned from this pool of IP addresses for clients in this user role.
Captive Portal Profile (optional)	This assigns a Captive Portal profile to this role. For more details about Captive Portal profiles, see Captive Portal Authentication on page 299.
Captive Portal Check for Accounting	This setting is enabled by default. If disabled, RADIUS accounting is done for an authenticated users irrespective of the captive-portal profile in the role of an authenticated user. If enabled, accounting is not done as long as the user's role has a captive portal profile on it. Accounting will start when Auth/XML-Add/CoA changes the role of an authenticated user to a role which doesn't have captive portal profile.
Max Sessions	This parameter configures the maximum number of sessions per user in this role. If the sessions reach the maximum value, any additional sessions from this user that are reaching the threshold are blocked till the session usage count for the user falls back below the configured limit. The default is 65535. You can configure any value between 0-65535.

To a delete a user role in the WebUI:

1. Navigate to the **Configuration > Security > Access Control > User Roles** page.
2. Click the **Delete** button against the role you want to delete.



You cannot delete a user-role that is referenced to profile or server derived role. Deleting a server referenced role will result in an error. Remove all references to the role and then perform the delete operation.

In the CLI

```
(host)(config) #user-role web-guest
   access-list session web-only position 1
```

Assigning User Roles

A client is assigned a user role by one of several methods. A role assigned by one method may take precedence over one assigned by a different method. The methods of assigning user roles are, from lowest to highest precedence:

1. The initial user role or VLAN for unauthenticated clients is configured in the AAA profile for a virtual AP (see Access Points (APs) on page 485).
2. The user role can be derived from user attributes upon the client's association with an AP (this is known as a user-derived role). You can configure rules that assign a user role to clients that match a certain set of criteria. For example, you can configure a rule to assign the role VoIP-Phone to any client that has a MAC address that starts with bytes *xx:yy:zz*.User-derivation rules are executed *before* client authentication.
3. The user role can be the default user role configured for an authentication method, such as 802.1x or VPN. For each authentication method, you can configure a default role for clients who are successfully authenticated using that method.
4. The user role can be derived from attributes returned by the authentication server and certain client attributes (this is known as a *server-derived role*). If the client is authenticated via an authentication server, the user role for the client can be based on one or more attributes returned by the server during authentication, or on client attributes such as SSID (even if the attribute is not returned by the server). Server-derivation rules are executed *after* client authentication.
5. The user role can be derived from Dell Vendor-Specific Attributes (VSA) for RADIUS server authentication. A role derived from a Dell VSA takes precedence over any other user roles.

The following sections describe the methods of assigning user roles.

Assigning User Roles in AAA Profiles

An AAA profile defines the user role for unauthenticated clients (initial role) as well as the default user role for MAC and 802.1x authentication. To configure user roles in the AAA profile:

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > AAA Profiles** page.
2. Select the default profile or a user-defined AAA profile.
3. Click the **Initial Role** drop-down list, and select the desired user role for unauthenticated users.
4. Click the **802.1x Authentication Default Role** drop-down list and select the desired user role for users who have completed 802.1x authentication.
5. Click the **MAC Authentication Default Role** drop-down list and select the desired user role for clients who have completed MAC authentication.
6. Click **Apply**.

In the CLI

```
(host)(config) #aaa profile <profile>
   initial-role <role>
   dot1x-default-role <role>
   mac-default-role <role>
```

For additional information on creating AAA profiles, see WLAN Authentication on page 421.

Working with User-Derived VLANs

Attributes derived from the client's association with an AP can be used to assign the client to a specific role or VLAN, as user-derivation rules are executed before the client is authenticated.

You configure the user role or VLAN to be assigned to the client by specifying condition rules; when a condition is met, the specified user role or VLAN is assigned to the client. You can specify more than one condition rule; the order of rules is important as the first matching condition is applied. You can optionally add a description of the user rule.

Table 63 describes the conditions for which you can specify a user role or VLAN.

Table 63: *Conditions for a User-Derived Roleor VLAN*

Rule Type	Condition	Value
BSSID: Assign client to a role or VLAN based upon the BSSID of AP to which client is associating.	One of the following: • contains • ends with • equals • does not equal • starts with	MAC address (xx:xx:xx:xx:xx:xx)
DHCP-Option: Assign client to a role or VLAN based upon the DHCP signature ID.	One of the following: • equals • starts with	DHCP signature ID. **NOTE:** This string is *not* case sensitive.
DHCP-Option-77: Assign client to a role or VLAN based upon the user class identifier returned by DHCP server.	equals	string
Encryption: Assign client to a role or VLAN based upon the encryption type used by the client.	One of the following: • equals • does not equal	• Open (no encryption) • WPA/WPA2 AES • WPA-TKIP (static or dynamic) • Dynamic WEP • WPA/WPA2 AES PSK • Static WEP • xSec
ESSID: Assign client to a role or VLAN based upon the ESSID to which the client is associated	One of the following: • contains • ends with • equals • does not equal • starts with • value of (does not take *string*; attribute value is used as role)	string
Location: Assign client to a role or VLAN based upon the ESSID to which the client is associated	One of the following: • equals • does not equal	string
MAC address of the client	One of the following: • contains • ends with • equals • does not equal • starts with	MAC address (xx:xx:xx:xx:xx:xx)

Understanding Device Identification

The device identification feature allows you to assign a user role or VLAN to a specific device type by identifying a DHCP option and signature for that device. If you create a user rule with the **DHCP-Option** rule type, the first two characters in the **Value** field must represent the hexadecimal value of the DHCP option that this rule should match, while the rest of the characters in the **Value** field indicate the DHCP signature the rule should match. To create a rule that matches DHCP option 12 (host name), the first two characters of the in the **Value** field must be the hexadecimal value of 12, which is 0C. To create a rule that matches DHCP option 55, the first two characters in the **Value** field must be the hexadecimal value of 55, which is 37.

The following table describes some of the DHCP options that are useful for assigning a user role or VLAN.

DHCP Option values

DHCP Option	Description	Hexadecimal Equivalent
12	Host name	0C
55	Parameter Request List	37
60	Vendor Class Identifier	3C
81	Client FQDN	51

The device identification features in ArubaOS can also automatically identify different client device types and operating systems by parsing the User-Agent strings in the client's HTTP packets. To enable this feature, select the **Device Type Classification** option in the AP's AAA profile. For details, see WLAN Authentication on page 421.

Configuring a User-derived VLAN in the WebUI

1. Navigate to the **Configuration > Security > Authentication > User Rules** page.
2. Click **Add** to add a new set of derivation rules. Enter a name for the set of rules, and click **Add**. The name appears in the User Rules Summary list.
3. In the User Rules Summary list, select the name of the rule set to configure rules.
4. Click **Add** to add a rule. For **Set Type**, select the VLAN name or ID from the **VLAN** the drop-down menu. (You can select **VLAN** to create d>erivation rules for setting the VLAN assigned to a client.)
5. Configure the condition for the rule by setting the Rule Type, Condition, Value parameters and optional description of the rule. See Table 63 for descriptions of these parameters.
6. Select the role assigned to the client when this condition is met.
7. Click **Add**.
8. You can configure additional rules for this rule set. When you have added rules to the set, use the up or down arrows in the Actions column to modify the order of the rules. (The first matching rule is applied.)
9. Click **Apply**.
10. (Optional) If the rule uses the DHCP-Option condition, best practices is to enable the Enforce DHCP parameter in the AP group's AAA profile, which requires users to complete a DHCP exchange to obtain an IP address. For details on configuring this parameter in an AAA profile, see WLAN Authentication on page 421.

Configuring a User-derived Role or VLAN in the CLI

```
(host)(config) #aaa derivation-rules user <name>
  set role|vlan
  condition bssid|dhcp-option|dhcp-option-77|encryption-type|essid|location|macaddr
  contains|ends-with|equals|not-equals|starts-with|value-of <string>
  set-value <role>
```

```
position <number>
```

See Table 63 for descriptions of these parameters.

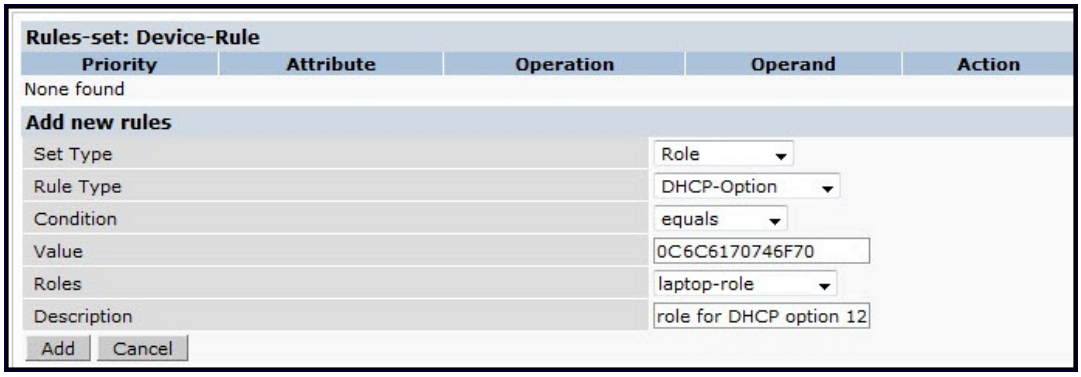
User-Derived Role Example

The example rule shown in Figure 37 below sets a user role for clients whose host name (DHCP option 12) has a value of 6C6170746F70, which is the hexadecimal equivalent of the ASCII string *laptop*. The first two digits in the **Value** field are the hexadecimal value of 12 (which is 0C), followed by the specific signature to be matched.



There are many online tools available for converting ASCII text to a hexadecimal string.

Figure 37 *DHCP Option Rule*



To identify DHCP strings used by an individual device, access the command-line interface in config mode and issue the following command to include DHCP option values for DHCP-DISCOVER and DHCP-REQUEST frames in the controller's log files:

```
logging level debugging network process dhcpd
```

Now, connect the device you want to identify to the network, and issue the CLI command `show log network`. The sample below is an example of the output that may be generated by this command.

```
(host)(config) #show log network all | include DISCOVER

Feb 26 02:50:34 :202534:  <DBUG> |dhcpdwrap| |dhcp| Datapath vlan1: DISCOVER 00:19:d2:01:0b:84 Options 74:01
3d:010019d2010b84 0c:736861626172657368612d39393730 3c:4d53465420352e30 37:010f03062c2e2f1f21f92b

Feb 26 02:50:42 :202534:  <DBUG> |dhcpdwrap| |dhcp| Datapath vlan1: DISCOVER 00:19:d2:01:0b:84 Options 74:01
3d:010019d2010b84 0c:736861626172657368612d39393730 3c:4d53465420352e30 37:010f03062c2e2f1f21f92b

Feb 26 02:50:42 :202534:  <DBUG> |dhcpdwrap| |dhcp| Datapath vlan1: DISCOVER 00:19:d2:01:0b:84 Options 74:01
3d:010019d2010b84 0c:736861626172657368612d39393730 3c:4d53465420352e30 37:010f03062c2e2f1f21f92b

Feb 26 02:53:03 :202534:  <DBUG> |dhcpdwrap| |dhcp| Datapath vlan10: DISCOVER 00:26:c6:52:6b:7c Options 74:01
3d:010026c6526b7c 0c:41525542412d46416c73653232 3c:4d53465420352e30 37:010f03062c2e2f1f21f92b 2b:dc00

...


(host) #show log network all| include REQUEST

Feb 26 02:53:04 :202536:  <DBUG> |dhcpdwrap| |dhcp| Datapath vlan10: REQUEST 00:26:c6:52:6b:7c reqIP=10.10.10.254
Options 3d:010026c6526b7c 36:0a0a0a02 0c:41525542412d46416c73653232
51:00000041525542412d46416c736532322e73757279612e636f6d 3c:4d53465420352e30 37:010f03062c2e2f1f21f92b 2b:dc0100

Feb 26 02:53:04 :202536:  <DBUG> |dhcpdwrap| |dhcp| Datapath vlan10: REQUEST 00:26:c6:52:6b:7c reqIP=10.10.10.254
```

Be aware that each device type may not have a unique DHCP fingerprint signature. For example, devices from different manufacturers may use vendor class identifiers that begin with similar strings. If you create a DHCP-Option rule that uses the **starts-with** condition instead of the **equals** condition, the rule may assign a role or VLAN to more than one device type.

RADIUS Override of User-Derived Roles

This feature introduces a new RADIUS vendor specific attribute (VSA) named "Aruba-No-DHCP-Fingerprint," value 14. This attribute signals the RADIUS Client (controller) to ignore the DHCP Fingerprint user role and VLAN change post L2 authentication. This feature applies to both CAP and RAP in tunnel mode and for the L2 authenticated role only.

Configuring a Default Role for Authentication Method

For each authentication method, you can configure a default role for clients who are successfully authenticated using that method. To configure a default role for an authentication method:

In the WebUI

1. Navigate to the **Configuration > Security > Authentication** page.
2. To configure the default user role for MAC or 802.1x authentication, select the **AAA Profiles** tab. Select the AAA profile. Enter the user role for MAC Authentication Default Role or 802.1x Authentication Default Role.
3. To configure the default user role for other authentication methods, select the **L2 Authentication** or **L3 Authentication** tab. Select the authentication type (Stateful 802.1x or stateful NTLM for L2 Authentication, Captive Portal or VPN for L3 Authentication), and then select the profile. Enter the user role for Default Role.
4. Click **Apply**.

For additional information on configuring captive portal authentication, see Captive Portal Authentication on page 299.

In the CLI

To configure the default user role for MAC or 802.1x authentication:

```
(host)(config) #aaa profile <profile>
  mac-default-role <role>
  dot1x-default-role <role>
```

To configure the default user role for other authentication methods:

```
(host)(config) #aaa authentication captive-portal <profile>
  d>efault-role <role>
(host)(config) #aaa authentication stateful-dot1x
  d>efault-role <role>
(host)(config) #aaa authentication stateful-ntlm
  d>efault-role <role>
(host)(config) #aaa authentication vpn
  d>efault-role <role>
```

Configuring a Server-Derived Role

If the client is authenticated through an authentication server, the user role for the client can be based on one or more attributes returned by the server during authentication. You configure the user role to be derived by specifying condition rules; when a condition is met, the specified user role is assigned to the client. You can specify more than one condition rule; the order of rules is important as the first matching condition is applied. You can also define server rules based on client attributes such as ESSID, BSSID, or MAC address, even though these attributes are not returned by the server.

For information about configuring a server-derived role, see Configuring Server-Derivation Rules on page 242.

Configuring a VSA-Derived Role

Many Network Address Server (NAS) vendors, including Dell, use VSAs to provide features not supported in standard RADIUS attributes. For Dell systems, VSAs can be employed to provide the user role and VLAN for

RADIUS-authenticated clients, however the VSAs must be present on your RADIUS server. This involves defining the vendor (Dell) and/or the vendor-specific code (14823), vendor-assigned attribute number, attribute format (such as string or integer), and attribute value in the RADIUS dictionary file. VSAs supported on controllers conform to the format recommended in RFC 2865, "Remote Authentication Dial In User Service (RADIUS)".

For more information on Dell VSAs, see RADIUS Server VSAs on page 228. Dictionary files that contain Dell VSAs are available on the Dell support website for various RADIUS servers. Log into the Dell support website to download a dictionary file from the Tools folder.

Understanding Global Firewall Parameters

Table 64 describes optional firewall parameters you can set on the controller for IPv4 traffic. To set these options in the WebUI, navigate to the **Configuration > Advanced Services > Stateful Firewall > Global Setting** page and select or enter values in the IPv4 column. To set these options in the CLI, use the **firewall** configuration commands.

See IPv6 Support on page 175 for information about configuring firewall parameters for IPv6 traffic.

Table 64: *IPv4 Firewall Parameters*

Parameter	Description
Monitor Ping Attack (per 30 seconds)	Number of ICMP pings per 30 second, which if exceeded, can indicate a denial of service attack. Valid range is 1-16384 pings per 30 seconds. Recommended value is 120. Default: No default
Monitor TCP SYN Attack rate (per 30 seconds)	Number of TCP SYN messages per 30 second, which if exceeded, can indicate a denial of service attack. Valid range is 1-16384 pings per 30 seconds. Recommended value is 960. Default: No default
Monitor IP Session Attack (per 30 seconds)	Number of TCP or UDP connection requests per 30 second, which if exceeded, can indicate a denial of service attack. Valid range is 1-16384 requests per 30 seconds. Recommended value is 960. Default: No default
Monitor/Police ARP Attack (non Gratuitous ARP) rate (per 30 seconds)	Number of ARP packets (other than Gratuitous ARP packets) per 30 seconds, which if exceeded, can indicate a denial of service attack. Valid range is 1-16384 packets per 30 seconds. Recommended value is 960. Default: No default **NOTE:** Blacklisting of wired clients is not supported.
Monitor/Police CP Attack rate (per 30 seconds)	Rate of misbehaving user's traffic, which if exceeded, can indicate a denial or service attack. Recommended value is 3000 frames per 30 seconds. Default: No default
Monitor/Police Gratuitous ARP Attack rate (per 30 seconds)	Number of Gratuitous ARP packets per 30 seconds, which if exceeded, can indicate denial of service attack. Valid range is 1-16384 packets per 30 seconds. Recommended value is 50. Default: 50 **NOTE:** Blacklisting of wired clients is not supported.

Parameter	Description
Deny Inter User Bridging	Prevents the forwarding of Layer-2 traffic between wired or wireless users. You can configure user role policies that prevent Layer-3 traffic between users or networks but this does not block Layer-2 traffic. This option can be used to prevent traffic, such as Appletalk or IPX, from being forwarded. Default: Disabled
Deny Inter User Traffic	Denies traffic between untrusted users by disallowing layer2 and layer3 traffic. This parameter does not depend on the deny-inter-user-bridging parameter being enabled or disabled. Default: Disabled
Deny Source Routing	Permits the firewall to reject and log packets with the specified IP options loose source routing, strict source routing, and record route. Note that network packets where the IPv6 source or destination address of the network packet is defined as an “link-local address (fe80::/64) are permitted. Default: Disabled
Deny All IP Fragments	Drops all IP fragments. **NOTE:** Do not enable this option unless instructed to do so by a Dell representative. Default: Disabled
Enforce TCP Handshake Before Allowing Data	Prevents data from passing between two clients until the three-way TCP handshake has been performed. This option should be disabled when you have mobile clients on the network as enabling this option will cause mobility to fail. You can enable this option if there are no mobile clients on the network. Default: Disabled
Prohibit IP Spoofing	Enables detection of IP spoofing (where an intruder sends messages using the IP address of a trusted client). When this option is enabled, source and destination IP and MAC addresses are checked for each ARP request/response. Traffic from a second MAC address using a specific IP address is denied, and the entry is not added to the user table. Possible IP spoofing attacks are logged and an SNMP trap is sent. Default: Enabled
Prohibit RST Replay Attack	When enabled, closes a TCP connection in both directions if a TCP RST is received from either direction. You should not enable this option unless instructed to do so by a Dell representative. Default: Disabled
Log ICMP Errors	Enables logging of received ICMP errors. You should not enable this option unless instructed to do so by a Dell representative. Default: Disabled
Stateful SIP Processing	Disables monitoring of exchanges between a voice over IP or voice over WLAN device and a SIP server. This option should be enabled only when there is no VoIP or VoWLAN traffic on the network. Default: Disabled (stateful SIP processing is enabled)
Allow Tri-session with DNAT	Allows three-way session when performing destination NAT. This option should be enabled when the controller is *not* the default gateway for wireless clients and the default gateway is behind the controller. This option is typically used for captive portal configuration. Default: Disabled.

Parameter	Description
Amsdu Configuration	Enables handling AMSDU traffic from clients. Default: Disabled
Session Mirror Destination	Destination (IP address or port) to which mirrored session packets are sent. This option is used only for troubleshooting or debugging. Packets can be mirrored in multiple ACLs, so only a single copy is mirrored if there is a match within more than one ACL. You can configure the following: • Ethertype to be mirrored with the Ethertype ACL mirror option. • IP flows to be mirrored with the session ACL mirror option. • MAC flows to be mirrored with the MAC ACL mirror option. • If you configure both an IP address and a port to receive mirrored packets, the IP address takes precedence. Default: N/A
Session Idle Timeout (sec)	Set the time, in seconds, that a non-TCP session can be idle before it is removed from the session table. Specify a value in the range 16-259 seconds. You should not set this option unless instructed to do so by a Dell representative. Default: 15 seconds
Disable FTP Server	Disables the FTP server on the controller. Enabling this option prevents FTP transfers. You should not enable this option unless instructed to do so by a Dell representative. Default: Disabled (FTP server is enabled)
GRE Call ID Processing	Creates a unique state for each PPTP tunnel. You should not enable this option unless instructed to do so by a Dell representative. Default: Disabled
Per-packet Logging	Enables logging of every packet if logging is enabled for the corresponding session rule. Normally, one event is logged per session. If you enable this option, each packet in the session is logged. You should not enable this option unless instructed to do so by a Dell representative, as doing so may create unnecessary overhead on the controller. Default: Disabled (per-session logging is performed)
Broadcast-filter ARP	Reduces the number of broadcast packets sent to VoIP clients, thereby improving the battery life of voice handsets. You can enable this option for voice handsets in conjunction with increasing the DTIM interval on clients. Default: Disabled
Prohibit ARP Spoofing	Detects and prohibits arp spoofing. When this option is enabled, possible arp spoofing attacks are logged and an SNMP trap is sent. Default: Disabled
Prevent DHCP exhaustion	Enable check for DHCP client hardware address against the packet source MAC address. This command checks the frame's source-MAC against the DHCPv4 client hardware address and drops the packet if it does not match. Enabling this feature prevents a client from submitting multiple DHCP requests with different hardware addresses, thereby preventing DHCP pool depletion. Default: Disabled

Parameter	Description
Session VOIP Timeout (sec)	Sets the idle session timeout for sessions that are marked as voice sessions. If no voice packet exchange occurs over a voice session for the specified time, the voice session is removed. Range is 16 – 300 seconds. Default: 300 seconds
Stateful H.323 Processing	Disables stateful H.323 processing. Default: Enabled
Stateful SCCP Processing	Disables stateful SCCP processing. Default: Disabled
Only allow local subnets in user table	Adds only IP addresses, which belong to a local subnet, to the user-table. Default: Disabled
Session mirror IPSEC	Configures session mirroring of all frames that are processed by IPsec. Frames are sent to IP address specified by the session-mirror-destination option. **NOTE:** Use this option for debugging or troubleshooting only. Default: Disabled
Session-tunnel FIB	Enable session-tunnel based forwarding. **NOTE:** Best practices is to enable this parameter only during maintenance window or off-peak production hours.
Multicast automatic shaping	Enables multicast optimization and provides excellent streaming quality regardless of the amount of VLANs or IP IGMP groups that are used. Default: Disabled
Stateful VOCERA Processing	Disables stateful VOCERA processing. Default: Disabled
Stateful UA Processing	Disables stateful UA processing. Default: Disabled
Enforce bw contracts for broadcast traffic	Applies bw contracts to local subnet broadcast traffic.
Enforce TCP Sequence numbers	Enforces the TCP sequence numbers for all packets. Default:Disabled
Enforce WMM Voice Priority Matches Flow Content	If traffic to or from the user is inconsistent with the associated QoS policy for voice, the traffic is reclassified to best effort and data path counters incremented. Default: Disabled
Rate limit CP untrusted ucast traffic (pps)	Specifies the untrusted unicast traffic rate limit. Range is 1-65535 pps. Default: 9765 pps
Rate limit CP untrusted mcast traffic (pps)	Specifies the untrusted multicast traffic rate limit. Range is 1-65535 pps. Default: 1953 pps
Rate limit CP trusted ucast traffic (pps)	Specifies the trusted unicast traffic rate limit. Range is 1-65535 pps. Default: 65535 pps
Rate limit CP trusted mcast traffic (pps)	Specifies the trusted multicast traffic rate limit. Range is 1-65535 pps. Default: 1953 pps

Parameter	Description
Rate limit CP route traffic (pps)	Specifies the traffic rate limit that needs ARP requests. Range is 1-65535 pps. Default: 976 pps
Rate limit CP session mirror traffic (pps)	Specifies the session mirrored traffic forwarded to the controller. Range is 1-65535 pps. Default: 976 pps
Rate limit CP auth process traffic (pps)	Specifies the traffic rate limit that is forwarded to the authentication process. Range is Range is 1-65535 pps. Default: 976 pps

Using AppRF 2.0

The AppRF 2.0 feature improves application visibility and control by allowing you to configure access control list (ACL) and bandwidth-control applications and application categories. AppRF 2.0 supports a Deep Packet Inspection (DPI) engine for application detection for over a thousand applications. All wired and wireless traffic that traverses the controller can now be categorized and controlled by application and application category.

AppRF 2.0 provides the ability to:

- permit or deny an application or application category for a specific role. For example, you can block bandwidth monopolizing applications on a guest role within an enterprise.
- rate limit an application or application category, such as video streaming applications, globally or for a specific role.
- mark different L2/L3 Quality of Service (QoS) for an application or application category for a user role. For example, you can mark video and voice sessions that originate from wireless users with different priorities so that traffic is prioritized accordingly in your network.

To configure AppRF 2.0, see the following topics:

- Enabling Deep Packet Inspection (DPI) on page 381
- Configuring Policies for AppRF 2.0 on page 382
- Configuring Bandwidth Contracts for AppRF 2.0 on page 384

Enabling Deep Packet Inspection (DPI)

For application and application category specific configuration to take affect, you must first enable DPI.

In the WebUI

1. Navigate to **Configuration > Advanced Services > Stateful Firewall > Global Settings**.
2. Check the **Enable Deep Packet Inspection** option. To disable DPI, uncheck the checkbox.
3. Click **Apply**.
4. Reload the controller.

In the CLI

To enable global DPI:

```
(host)(config) #firewall dpi
(host) #reload
```



You must reboot (reload) the controller after you enable DPI for global classification to take effect.

To disable global DPI:

```
(host)(config)# no firewall dpi
```



You must reboot (reload) the controller after you disable DPI for global classification to be disabled.

Show Command Output

The **show datapath session** output now includes:

- A new parameter, **show datapath session dpi**, displays application ID, application name, and the following ACL/ACE index information for a given session:
  - **AclVersion:** This is used to store the current version number of the ACL that is used at session creation time and is used for troubleshooting purposes.
  - **PktsDpi:** The number of packets sent to the DPI engine for a given session.
  - **AceIdx:** The Index of the Access List entry (in a given ACL) that triggered a match during session creation.
  - **DpiTIdx:** This is an index to the DPI engine Tbl and is only used for troubleshooting purposes.
- A new flag, **A - Application Firewall Inspect**, indicates that a flow is being subjected to DPI.

Configuring Policies for AppRF 2.0

Access control lists now contain new application and application category options that let you permit or deny an application or application category on a given role. See the Dashboard Monitoring AppRF topic for details about configuring policies from the Dashboard.

How ACL Works with AppRF

A session entry proceeds through two phases: the application detection phase (phase1) and the post-application detection phase (phase 2). A session ACL is applied in phase1 and in phase 2.

In phase1, if the session ACL lookup results in an L3/L4 ACE entry request, the traffic pertaining to the session is guided by this L3/L4 ACE entry. However, if the session ACL lookup results in an application/application category specific ACE entry, the enforcement is postponed until phase 2. Once the application is determined, the session ACL is re-applied with "application/application category" information to determine the final action on the traffic.

Global Session ACL

The Global Session ACL is used to configure ACL rules that span across or are common to all roles. They are applied to all roles. The "global-sacl" rules take precedence over any other ACLs that may be in the user role.

A new session ACL has been added named "global-sacl." This session, by default, is in position one for every user role configured on the controller. The global-sacl session ACL has the following properties:

- It cannot be deleted.
- It always remains at position one in every role and its position cannot be modified.
- It contains only application rules.
- It can be modified in the WebUI, CLI, and dashboard on a master controller.
- Any modifications to it resulst in the regeneration of ACE's of all roles.

Role Default Session ACL

You can configure role-specific application configuration using the WebUI and dashboard. For example, you can deny the facebook application on the guest role using the CLI or dashboard without having to change the

firewall configuration. This per-user role configuration from WebUI or Dashboard is placed in the Role Default Session ACL.

A new role session ACL named apprf-“role-name”-sacl has been added. This session, by default, is in position two for every user role configured on the controller.

The string "apprf" is added to the beginning and "sacl" to the end of a role’s name to form a controllerunique name for role default session ACL. This session ACL is in position two of the given user role after the global session ACL and takes the next higher priority after global policy rules.

The predefined role session ACL has the following properties:

- It cannot be deleted through the WebUI or CLI. It it is only deleted automatically when the corresponding role is deleted.
- It always remains at position 2 in every role and its position cannot be modified.
- It contains only application rules.
- It can be modified using the WebUI, CLI, or dashboard on a master controller, however any modification results in the regeneration of ACE’s for that role.
- It cannot be applied to any other role.

Session ACL Examples

The following CLI configuration shows how pre-classification and post-classification occurs during enforcement.

```
Any any app skype permit
Any any deny
```

Each application has an implicit set of ports that are used for communication. In phase 1, if an application ACE entry is hit, the traffic matching this application’s implicit port is allowed (as governed by the application ACE). The DPI engine can monitor the exchange on these ports and determine the application. Once the application is determined, phase 2 occurs when an evaluation is done to determine the final outcome for the session.

The following CLI configuration example is a user role with both the global and role session ACLs:

User-role employee

```
ip access-list session global-sacl
ip access-list session apprf-employee-sacl
ip access-list session control
   any any app gmail-chat permit
   any any app youtube permit
   any any any deny
```

This example shows a DPI rule along with a L3/L4 rule with forwarding action in the same ACL.

```
ip access-list session AppRules
   any any app Facebook permit tos 45
   any any app YouTube deny
   any any appcategory peer-to-peer deny
   any any tcp 23 permit
   network 40.1.0.0/16 any tcp 80 permit tos 60
   network 20.1.0.0/16 any tcp 80 src-nat
!
ip access-list session NetRules
   network 80.0.0.0/24 any tcp 80 deny
   network 60.0.0.0/24 any tcp 80 dual-nat pool <pool1>
   network 10.0.0.0/24 any tcp 80 dst-nat
!
user-role Role1
   session-acl AppRules
   session-acl NetRules
!
```

In the WebUI

1. Navigate to **Configuration > Access Control > Policies.**
2. Click **Add/Edit**.
3. Click **Add** under Rules/IP Version.
4. Select **application** or **application category** from the **Service** drop-down menu and select configuration options.
5. Click **Apply**.

In the CLI

To configure the ACL application-specific parameters using the command-line interface, access the command-line interface in config mode, run the following commands:

```
ip access-list session <accname>
<source> <dest> <service> <action>[<extended action>]
```

Configuring Bandwidth Contracts for AppRF 2.0

Bandwidth contract configuration lets you configure bandwidth contracts for both the global or application-specific levels.

Global Bandwidth Contract Configuration

You can configure bandwidth contracts to limit application and application categories on an application or global level.

In the CLI

To configure global bandwidth contracts:

```
(host)(config) #dpi global-bandwidth-contract[app|appcategory]
<name>[downstream|upstream][kbits|mbits]<256..2000000>
```

To show global bandwidth contract configuration output:

```
(host) #show dpi global-bandwidth-contract all
(host) #show dpi global-bandwidth-contract app name
(host) #show dpi global-bandwidth-contract appcategory name
```

Role-Specific Bandwidth Contracts

Application-specific bandwidth contracts (unlike "generic" bandwidth-contracts) allow you to control or reserve rates for specific applications only on a per-role basis. An optional exclude list is provided that allows you to exclude applications or application categories on which a generic user/role bandwidth-contract is not applied.

Using an Exclude List

The exclude list enables you to give specific enterprise mission-critical applications priority over other user traffic. An enterprise may have well known applications such as Microsoft Exchange, SAP, Oracle, accounting and finance applications, and other enterprise resource planning (ERP) or customer relationship management (CRM) applications.

An exclude list helps you to prioritize these mission-critical enterprise applications.

Instead of enumerating bandwidth limits for each application individually on a per-user/per-role basis, you can configure a single bandwidth contract on a per-user/per-role to limit all non-mission critical applications. You can then exclude all mission-critical applications by placing them in an exclude list. This way all mission-critical applications will not be rate-limited.

Important points regarding bandwidth contracts include:

- Application bandwidth contracts are per-role by default.

- When an application bandwidth-contract is configured for both a cateogry and an application within the category, always apply the most specific bandwidth contract.

In the WebUI

1. Navigate to **Configuration > Security > Access Control > User Roles**.
2. Click **Add** to create a new user role or **Edit** to modify an existing role.
3. Select the **Bandwidth Contracts** tab.
4. To exclude an application or application category, click the **Add** button below the **Exceptions** section, select an item from the **Name** drop-down menu and click **Done**.
5. To add an application or application category to a bandwidth contract, click **Add** under **Application Bandwidth Contracts**.
6. Select the application from the **Name** drop-down me and whether it is enforced.
7. Enter a name of the new bandwidth contract, the bandwidth in kpbs or mbps, and if downstream is enforced.
8. Select an option from the **Downstream** drop-down menu and Per Role, Per User, or Per AP Group from the adjacent drop-down menu.
9. Select additional configuration parameters from the **Misc. Configuration** pane.



Make sure that the **Enable Deep Packet Inspection** option is checked.

10. Click **Apply**.

In the CLI

To configure the bandwidth application-specific parameters using the CLI, access the command-line interface in config mode, and issue the following commands:

```
(host)config t #user-role <string>
(host)(config-role) #bw-contract exclude[app|app-cat]
<string>|[upstream|downstream]
```

To exclude a specific application or application category, use the following command.

```
(host)(config-role)#bw-contract exclude[app|app-cat]|<name>
```

ArubaOS and ClearPass Policy Manager (CPPM) include support for centralized policy definition and distribution. ArubaOS now supports downloadable roles. By using this feature, when CPPM successfully authenticates a user, the user is assigned a role by CPPM and if the role is not defined on the controller, the role attributes can also be automatically downloaded.

This chapter contains the following sections:

- Introduction on page 386
- Important Points to Remember on page 386
- Enabling Downloadable Role on a Controller on page 387
- Sample Configuration on page 387

Introduction

In order to provide highly granular per-user level access, user roles can be created when a user has been successfully authenticated. During the configuration of a policy enforcement profile at CPPM, the administrator can define a role that should be assigned to the user after successful authentication. In RADIUS authentication, when CPPM successfully authenticates a user, the user is assigned a role by CPPM and if the role is not defined on the controller, the role attributes can also be automatically downloaded. This feature supports roles obtained by the following authentication methods:

- 802.1x (wireless and wired users)
- MAC authentication
- Captive Portal

Important Points to Remember

- Under **Advanced** mode, CPPM does not perform any error checking to confirm accuracy of the role definition. Therefore, it is recommended that you review the role defined in CPPM prior to enabling this feature.
- Attributes that are listed below, herein referred to as whitelist role attributes, can be defined in CPPM.
  - **netdestination**
  - **netservice**
  - **ip access-list eth**
  - **ip access-list mac**
  - **ip access-list session**
  - **user-role**
- The above attributes that are referred to by a role definition must either be defined within the role definition itself or configured on the controller before the policy is downloaded.
- In CPPM, two or more attributes (as listed above) should not have the same name. Example below is considered invalid as both the attributes have **test** as the profile/net destination name.

```
qos-profile test
netdestination test
```

- An instance name (name of a whitelist role attribute as stated above) is case-sensitive. Attributes must adhere to the following rules:
  - Should not match any CLI option nested under a command from the whitelist.
  - Should not contain a number or a combination of numbers.
  - Should not contain any periods '.'.
  - Should not contain any spaces.

  Example below is considered as invalid configurations and will fail CPPM role download on a controller:

  ```
  netservice 'tcp' tcp 443
  ```

  The first instance of **tcp** is a user-defined field while the second is an operator of the **netservice** command. This violates the first rule.

  ```
  netdestination 'alias'
  ```

  The user-defined name **alias** is also a valid operator of the **netdestination** command. This violates the first rule.

  ```
  netdestination '10.1.5'
  ```

  This user-defined name uses both numbers and periods. This violates the second and third rule.

  ```
  ip access-list stateless '100'
  ```

  This user-defined name uses numbers. This violates the second rule.

  ```
  qos-profile emp role
  ```

  This profile name **emp role** contains spaces. This violates the fourth rule.

It is recommended that some naming convention similar to the CamelCase (mixture of upper and lower case letters in a single word) be used to avoid collisions with the CLI options in the role description.

Enabling Downloadable Role on a Controller

You can enable role download using the CLI or WebUI.

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > AAA Profiles**.
2. Select an AAA profile.
3. Check the **Download Role from CPPM** check box to enable role download.

Using the CLI

```
(host) (config) #aaa profile <profile-name>
(host) (AAA profile) #download-role
```

Sample Configuration

The following example shows the configuration details to integrate CPPM server with a controller to automatically download roles.

CPPM Server Configuration

Adding a Device

1. From the **Configuration > Network > Devices** page, click the **Add Device** link.
2. On the **Device** tab, enter the **Name**, **IP or Subnet Address**, and **RADIUS Shared Secret** fields. Keep the rest of the fields as default.

3. Click **Add**.

The fields are described in Figure 38 and Table 65.

Figure 38 *Device Tab*

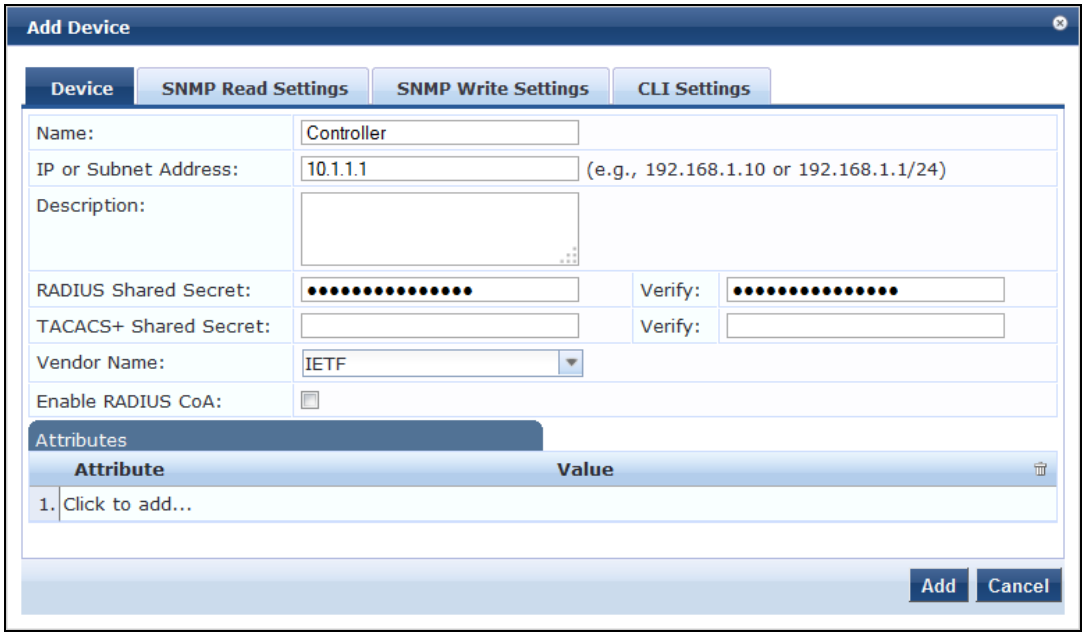


Table 65: *Device Tab*

Container	Description
Name	Specify the name or identity of the device.
IP or Subnet Address	Specify the IP address or subnet (example 10.1.1.1/24) of the device.
RADIUS Shared Secret	Enter and confirm a Shared Secret for each of the two supported request protocols.

Adding Enforcement Profile

1. From **Configuration > Enforcement > Profiles** page, click **Add Enforcement Profile**.
2. On the **Profile** tab, select **Aruba Downloadable Role Enforcement** from the **Template** drop-down list.
3. Enter the **Name** of the enforcement profile.
4. From the **Role Configuration Mode**, select **Advanced**.

   Keep the rest of the fields as default.
5. Click **Next**.

   For the rest of the configuration, see Advanced Role Configuration Mode.

The fields are described in Figure 39 and Table 66.

Figure 39 *Enforcement Profiles Page*

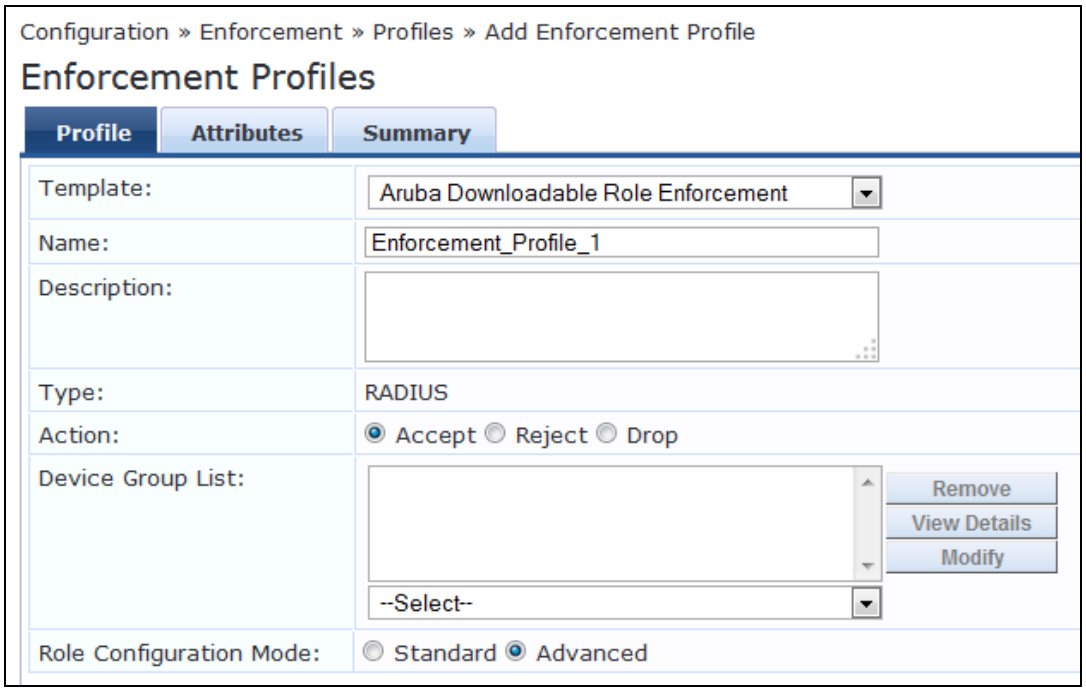


Table 66: *Enforcement Profiles Page*

Container	Description
Template	Policy Manager comes pre-packaged with several enforcement profile templates. In this example, select **Aruba Downloadable Role Enforcement** - RADIUS template that can be filled with user role definition to create roles that can be assigned to users after successful authentication.
Name	Specify the name of the enforcement profile.
Role Configuration Mode	Standard—Configure enforcement profile role using standard mode. Advanced—Configure enforcement profile role using advanced mode.

Advanced Role Configuration Mode

1. On the **Attributes** tab, select **Radius:Aruba** from the **Type** drop-down list.
2. From the **Name** drop-down list, select **Aruba-CPPM-Role**.
3. In the **Value** field, enter the attribute for the downloadable-role.
4. Click the save icon to save the attribute.
5. Click **Save** to save the enforcement profile.

The fields are described in Figure 40 and Table 67.

Figure 40 *Enforcement Profiles Attributes Tab*

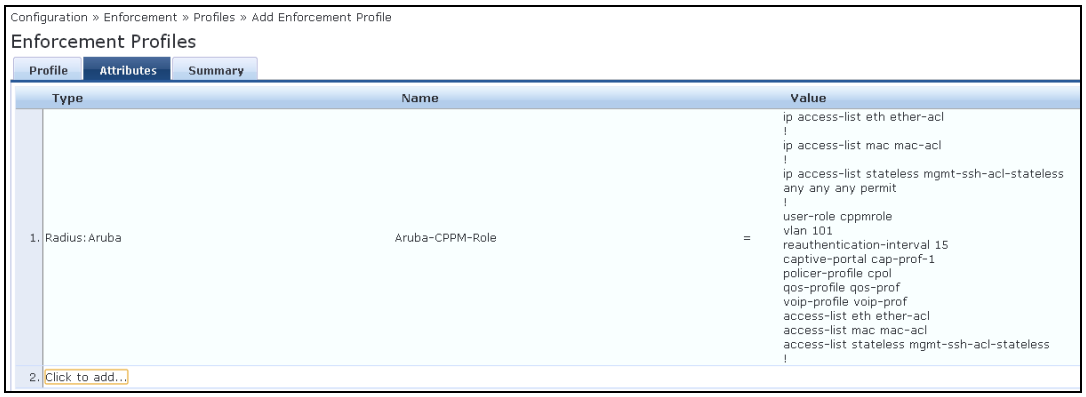


Table 67: *Enforcement Profiles Attributes Tab*

Container	Description
Type	Type is any RADIUS vendor dictionary that is pre-packaged with Policy Manager, or imported by the Administrator. This field is pre-populated with the dictionary names.
Name	Name is the name of the attribute from the dictionary selected in the Type field. The attribute names are pre-populated from the dictionary.
Value	Value is attribute for the downloadable role. You can enter free-form text to define the role and policy. **NOTE:** The maximum limit for free form text is 16,000 bytes.

Adding Enforcement Policy

1. From **Configuration > Enforcement > Policies** page, click **Add Enforcement Policy**.
2. On the **Enforcement** tab, enter the name of the enforcement policy.
3. From the **Default Profile** drop-down list, select **[Deny Access Profile]**.
   Keep the rest of the fields as default.
4. Click **Next**.

The fields are described in Figure 41 and Table 68.

Figure 41 *Enforcement Policies Enforcement Tab*

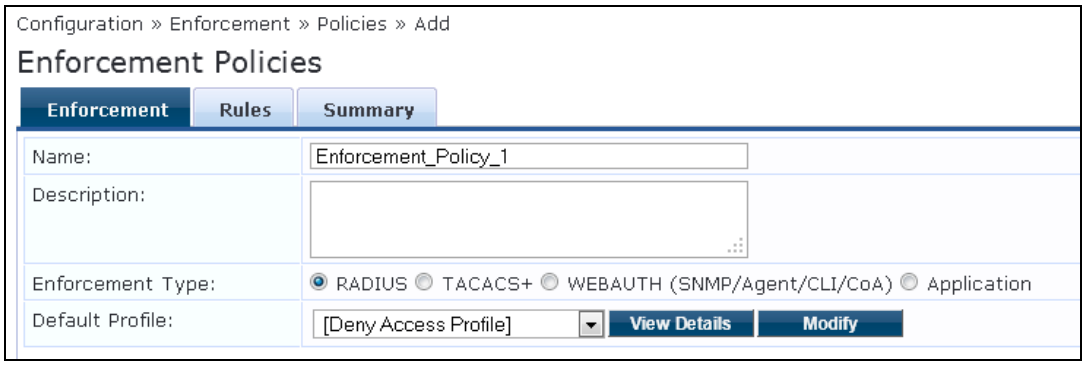


Table 68: *Enforcement Policies Enforcement Tab*

Container	Description
Name	Specify the name of the enforcement policy.
Default Profile	An Enforcement Policy applies Conditions (roles, health, and time attributes) against specific values associated with those attributes to determine the Enforcement Profile. If none of the rules matches, Policy Manager applies the Default Profile. See Adding Enforcement Profile on page 388 to add a new profile.

5. On the **Rules** tab, click **Add Rule**.
6. On the **Rules Editor** pop-up, select the appropriate values in the **Conditions** section and click the save icon.
7. In the **Enforcement Profiles** section, select the RADIUS enforcement profile that you created in step Adding Enforcement Profile on page 388 from the **Profile Names** drop-down list.
8. Click **Save**.

The fields are described in Figure 42 and Table 69.

Figure 42 *Enforcement Policies Rules Editor*

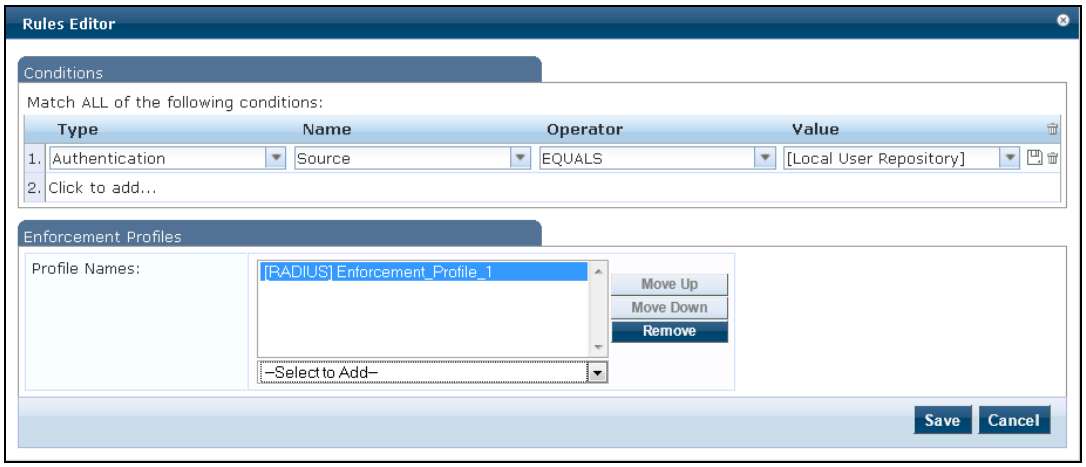


Table 69: *Enforcement Policies Rules Editor*

Container	Description
Type	The rules editor appears throughout the Policy Manager interface. It exposes different namespace dictionaries depending on Service type. When working with service rules, you can select **Authentication** namespace dictionary
Name	Drop-down list of attributes present in the selected namespace. In this example, select **Source**.
Operator	Drop-down list of context-appropriate (with respect to the attribute) operators. In this example, select **EQUALS**.
Value	Drop-down list of the Authentication source database. In this example, select **[Local User Repository]**.
Profile Names	Name of the RADIUS enforcement profile.

Adding Services

1. From the **Configuration > Services** page, click the **Add Service** link.
2. On the **Service** tab, select **802.1X Wired** from the **Type** drop-down-list.
3. In the **Name** field, enter the name of the service.

   Keep the rest of the fields as default.
4. Click **Next**.

The fields are described in Figure 43 and Table 70.

Figure 43 *Service Tab*

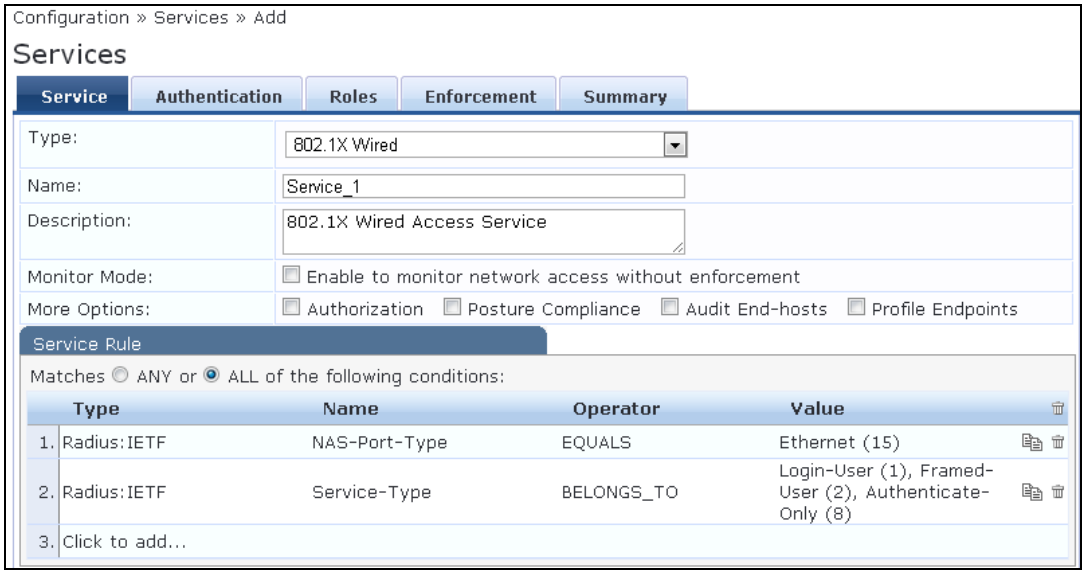


Table 70: *Service Tab*

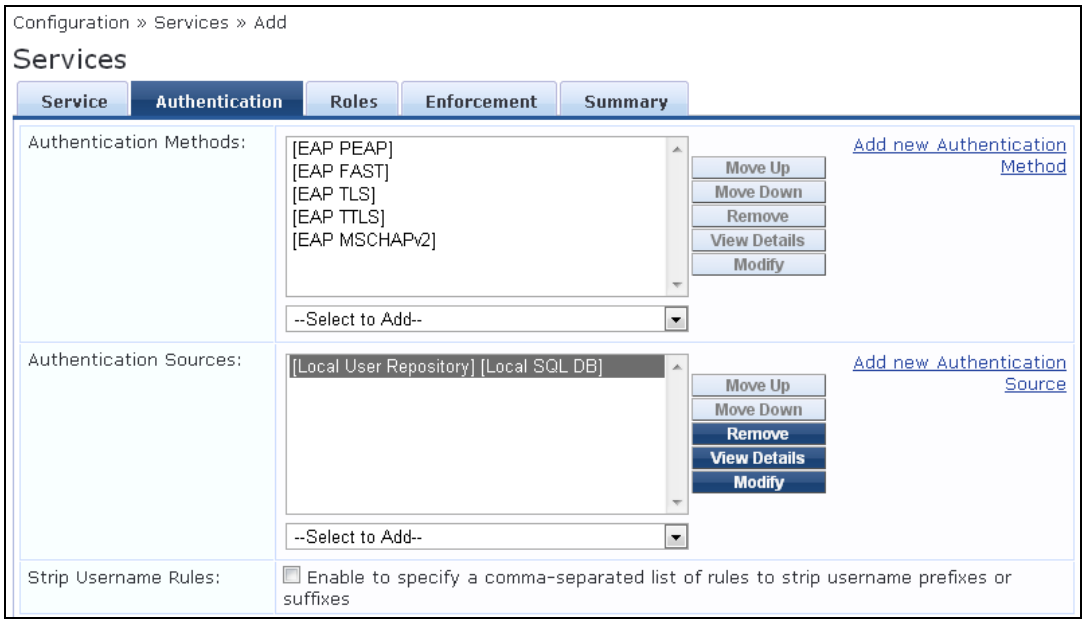
Container	Description
Type	Select the desired service type from the drop down menu. In this example, select **802.1X Wired**.
Name	Specify the name of the service.

5. On the **Authentication** tab, select **[Local User Repository] [Local SQL DB]** from the **Authentication Sources** drop-down list.

   Keep the rest of the fields as default.
6. Click **Next** twice.

The fields are displayed in Figure 44.

Figure 44 *Authentication Tab*

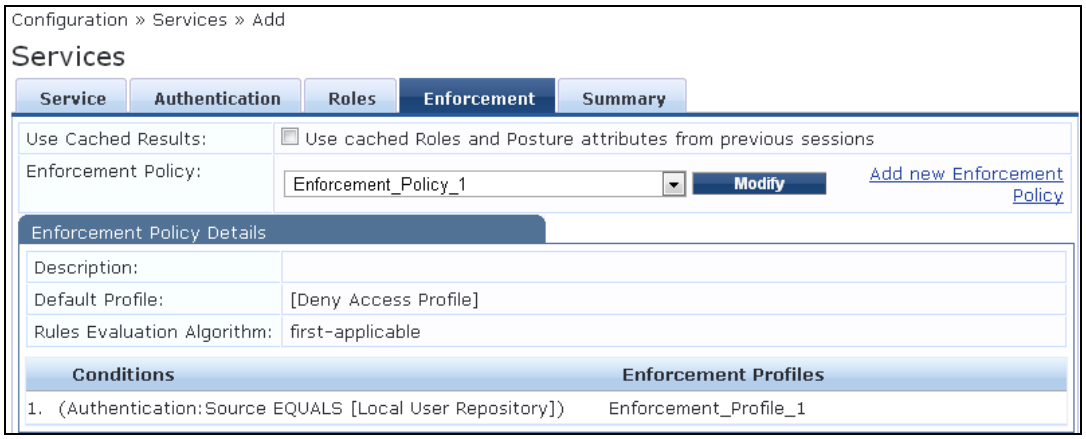


7. On the **Enforcement** tab, select the enforcement policy that you created in step Adding Enforcement Policy on page 390 from the **Enforcement Policy** drop-down list.

   Keep the rest of the fields as default.
8. Click **Save**.

The fields are displayed in Figure 45.

Figure 45 *Enforcement Tab*



For more configuration details on CPPM, see the *ClearPass Policy Manager 6.2 User Guide.*

Controller Configuration

Configuring CPPM Server on Controller

```
(host) (config) #aaa authentication-server radius cppm_server
(host) (RADIUS Server "cppm_server") #host <ip_address_of_cppm_server>
(host) (RADIUS Server "cppm_server") #key <shared_secret>
```

Configuring Server Group to include CPPM Server

```
(host) (config) #aaa server-group cppm_grp
(host) (Server Group "cppm_grp") #auth-server cppm_server
```

Configuring 802.1X Profile

```
(host) (config) #aaa authentication dot1x cppm_dot1x_prof
```

Configuring AAA Profile

```
(host) (config) #aaa profile cppm_aaa_prof
(host) (AAA Profile "cppm_aaa_prof") #authentication-dot1x cppm_dot1x_prof
(host) (AAA Profile "cppm_aaa_prof") #dot1x-server-group cppm_grp
(host) (AAA Profile "cppm_aaa_prof") #download-role
```

Show AAA Profile

```
(host) #show aaa profile cppm_aaa_prof
```

APs advertise WLANs to wireless clients by sending out beacons and probe responses that contain the WLAN's SSID and supported authentication and data rates. When a wireless client associates to an AP, it sends traffic to the AP's Basic Service Set Identifier (BSSID) which is usually the AP's MAC address.

In the Dell network, an AP uses a unique BSSID for each WLAN. Thus, a physical AP can support multiple WLANs. The WLAN configuration applied to a BSSID on an AP is called a *virtual AP*. You can configure and apply multiple virtual APs to an AP group or to an individual AP by defining one or more virtual AP profiles.

This chapter describes the following topics:

- Virtual AP Profiles on page 395
- Virtual AP Configuration Workflow on page 403
- Radio Resource Management (802.11k) on page 404
- BSS Transition Management (802.11v) on page 411
- Fast BSS Transition (802.11r) on page 412
- SSID Profiles on page 414
- WLAN Authentication on page 421
- High-Throughput Virtual APs on page 424
- Guest WLANs on page 428

Virtual AP Profiles

You can configure virtual AP profiles to provide different network access or services to users on the same physical network. For example, you can configure a WLAN to provide access to guest users and another WLAN to provide access to employee users through the same APs. You can also configure a WLAN that offers open authentication and Captive Portal access with data rates of 1 and 2 Mbps, and another WLAN that requires WPA authentication with data rates of up to 11 Mbps. You can apply both virtual AP configurations to the same AP or an AP group .

As an example, suppose there are users in both Edmonton and Toronto that access the same "Corpnet" WLAN. If the WLAN required authentication to an external server, users who associate with the APs in Toronto would want to authenticate with their local servers. In this case, you can configure two virtual APs that each reference a slightly different AAA profile; one AAA profile that references authentication servers in Edmonton and the other that references servers in Toronto (see Table 71).

Table 71: *Applying WLAN Profiles to AP Groups*

WLAN Profiles	"default" AP Group	"Toronto" AP Group
Virtual AP	"Corpnet-Ed"	"Corpnet-Tr"
SSID	"Corpnet"	"Corpnet"
AAA	"E-Servers"	"T-Servers"

You can apply multiple virtual AP profiles to individual APs. You can also apply the same virtual AP profile to one or more AP groups.

Configuring the Virtual AP Profile

Follow the procedures below to configure a Virtual AP profile using the WebUI or command-line interfaces.

Creating and Configuring a Profile

1. Navigate to **Configuration > Advanced Services > All Profiles**.
2. In the **Profiles** pane, expand the **Wireless LAN** menu.
3. Select **Virtual AP**. The list of existing Virtual AP profiles appears in the **Profile Details** pane.
4. Select the virtual AP profile you want to configure:
   - To configure an existing Virtual AP profile, select the name of the profile in the **Profile Details** pane.
   - To create a new Virtual AP profile, enter a name for the profile in the entry blank at the bottom of the **Profile Details** pane, then click **Add**. Select the name of the profile in the **Profile Details** pane.



Whenever you create a new virtual AP profile in the WebUI, the profile automatically contains the "default" SSID profile with the default "Dell-ap" ESSID. You must configure a new ESSID and SSID profile for the virtual AP profile before you apply the profile.

5. Configure the profile parameters described in Table 72.

   The Virtual AP profile is divided into two tabs, **Basic** and **Advanced**. The **Basic** tab displays only those configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. If you change a setting on one tab then click and display the other tab without saving your configuration, that setting will revert to its previous value.

Table 72: *Virtual AP Profile Parameters*

Parameter	Description
Basic Configuration Settings	
Virtual AP enable	Select the **Virtual AP enable** checkbox to enable or disable the virtual AP.
VLAN	The VLAN(s) into which users are placed in order to obtain an IP address. Click the drop-down list to select a configured VLAN, the click the arrow button to associate that VLAN with the virtual AP profile.
Forward mode	This parameter controls whether data is tunneled to the controller using generic routing encapsulation (GRE), bridged into the local Ethernet LAN (for remote APs), or a combination thereof depending on the destination (corporate traffic goes to the controller, and Internet access remains local). All forwarding modes support band steering, TSPEC/TCLAS enforcement, 802.11k and station blacklisting. Click the drop-down list to select one of the following forward modes: • **Tunnel:** The AP handles all 802.11 association requests and responses, but sends all 802.11 data packets, action frames and EAPOL frames over a GRE tunnel to the controller for processing. The controller removes or adds the GRE headers, decrypts or encrypts 802.11 frames and applies firewall rules to the user traffic as usual. Both remote and campus APs can be configured in tunnel mode. • **Bridge:** 802.11 frames are bridged into the local Ethernet LAN. When a remote AP or campus AP is in bridge mode, the AP (and not the controller) handles all 802.11 association requests and responses, encryption/decryption processes, and firewall enforcement. The 802.11e and 802.11k action frames are also

Parameter	Description
	processed by the AP, which then sends out responses as needed. An AP in bridge mode does not support captive portal authentication. Both remote and campus APs can be configured in bridge mode. Note that you must enable the control plane security feature on the controller before you configure campus APs in bridge mode. • **Split-Tunnel:** 802.11 frames are either tunneled or bridged, depending on the destination (corporate traffic goes to the controller, and Internet access remains local). A remote AP in split-tunnel forwarding mode handles all 802.11 association requests and responses, encryption/decryption, and firewall enforcement. the 802.11e and 802.11k action frames are also processed by the remote AP, which then sends out responses as needed. • **Decrypt-Tunnel:** Both remote and campus APs can be configured in decrypt-tunnel mode. When an AP uses decrypt-tunnel forwarding mode, that AP decrypts and decapsulates all 802.11 frames from a client and sends the 802.3 frames through the GRE tunnel to the controller, which then applies firewall policies to the user traffic. When the controller sends traffic to a client, the controller sends 802.3 traffic through the GRE tunnel to the AP, which then converts it to encrypted 802.11 and forwards to the client. This forwarding mode allows a network to utilize the encryption/decryption capacity of the AP while reducing the demand for processing resources on the controller. APs in decrypt-tunnel forwarding mode also manage all 802.11 association requests and responses, and process all 802.11e and 802.11k action frames. APs using decrypt-tunnel mode do have some limitations that not present for APs in regular tunnel forwarding mode. You must enable the control plane security feature on the controller before you configure campus APs in decrypt-tunnel forward mode. **NOTE:** Virtual APs in bridge or split-tunnel mode using static WEP should use key slots 2-4 on the controller. Key slot 1 should only be used with Virtual APs in tunnel mode.
Allowed band	The band(s) on which to use the virtual AP: • **a**—802.11a band only (5 GHz). • **g**—802.11b/g band only (2.4 GHz). • **all**—both 802.11a and 802.11b/g bands (5 GHz and 2.4 GHz). This is the default setting.
Band Steering	ARM's band steering feature encourages dual-band capable clients to stay on the 5GHz band on dual-band APs. This frees up resources on the 2.4GHz band for single band clients like VoIP phones. Band steering reduces co-channel interference and increases available bandwidth for dual-band clients, because there are more channels on the 5GHz band than on the 2.4GHz band. Dual-band 802.11n-capable clients may see even greater bandwidth improvements, because the band steering feature will automatically select between 40MHz or 20MHz channels in 802.11n networks. This feature is disabled by default, and must be enabled in a Virtual AP profile. The band steering feature supports both campus APs and remote APs that have a virtual AP profile set to **tunnel**, **split-tunnel** or **bridge** forwarding mode. Note, however, that if a campus or remote APs has virtual AP profiles configured in bridge or split-tunnel forwarding mode but no virtual AP in tunnel mode, those APs will gather information about 5G-capable clients independently and will not exchange this information with other APs that also have bridge or split-tunnel virtual APs only.
Steering Mode	Band steering supports the following three different band steering modes. • **Force-5GHz:** When the AP is configured in **force-5GHz** band steering mode, the AP will try to force 5Ghz-capable APs to use that radio band. • **Prefer-5GHz** (Default): If you configure the AP to use **prefer-5GHz** band

Parameter	Description
	steering mode, the AP will try to steer the client to 5G band (if the client is 5G capable) but will let the client connect on the 2.4G band if the client persists in 2.4G association attempts. • **Balance-bands:** In this band steering mode, the AP tries to balance the clients across the two radios in order to best utilize the available 2.4G bandwidth. This feature takes into account the fact that the 5Ghz band has more channels than the 2.4 Ghz band, and that the 5Ghz channels operate in 40MHz while the 2.5Ghz band operates in 20MHz.
Dynamic Multicast Optimization (DMO)	Enable/Disable dynamic multicast optimization. This parameter is disabled by default, and cannot be enabled without the PEFNG license.
Drop Broadcast and Multicast	Select the **Drop Broadcast and Multicast** checkbox to filter out broadcast and multicast traffic in the air. Do not enable this option for virtual APs configured in bridge forwarding mode. This configuration parameter is only intended for use for virtual APs in tunnel mode. In tunnel mode, all packets travel to the controller, so the controller is able to drop all broadcast traffic. When a virtual AP is configured to use bridge forwarding mode, most data traffic stays local to the AP, and the controller is not able to filter out that broadcast traffic. **IMPORTANT:** If you enable this option, you must also enable the **Broadcast-Filter ARP** parameter on the virtual AP profile to prevent ARP requests from being dropped. You can enable this parameter by checking the **Convert Broadcast ARP requests to unicast** check box as described in the following parameter description.
Convert Broadcast ARP requests to unicast	If enabled, all broadcast ARP requests are converted to unicast and sent directly to the client. You can check the status of this option using the **show ap active** and the **show datapath tunnel** command. If enabled, the output will display the letter **a** in the flags column. This configuration parameter is only intended for use for virtual APs in tunnel mode. In tunnel mode, all packets travel to the controller, so the controller is able to convert ARP requests directed to the broadcast address into unicast. When a virtual AP is configured to use bridge forwarding mode, most data traffic stays local to the AP, and the controller is not able to convert that broadcast traffic. Beginning with ArubaOS 6.1.3.2, this parameter is enabled by default. Behaviors associated with these settings are enabled upon upgrade to ArubaOS 6.1.3.2. If your controller supports clients behind a wireless bridge or virtual clients on VMware devices, you must disable this setting to allow those clients to obtain an IP address. In previous releases of ArubaOS, the virtual AP profile included two unique broadcast filter parameters; the **drop broadcast and multicast** parameter, which filtered out all broadcast and multicast traffic in the air except DHCP response frames (these were converted to unicast frames and sent to the corresponding client) and the **conert ARP requests to unicast** parameter, which converted broadcast ARP requests to unicast messages sent directly to the client. Starting with ArubaOS 6.1.3.2, the **Convert Broadcast ARP requests to unicast** setting includes the additional functionality of broadcast-filter all parameter, where DHCP response frames are sent as unicast to the corresponding client. This can impact DHCP discover/requested packets for clients behind a wireless bridge and virtual clients on VMware devices. Disable this option to resolve this issue and allow clients behind a wireless bridge or VMware devices to receive an IP address. Default: Enabled
Advanced Configuration Settings	
Dynamic Multicast Optimization (DMO) Threshold	Maximum number of high-throughput stations in a multicast group beyond which dynamic multicast optimization stops. Range: 2-255 stations Default: 6 stations.

Parameter	Description
Blacklist Time	Number of seconds that a client is quarantined from the network after being blacklisted. Default: 3600 seconds (1 hour)
Authentication Failure Blacklist Time	Time, in seconds, a client is blocked if it fails repeated authentication. The default setting is 3600 seconds (1 hour). A value of 0 blocks the client indefinitely.
Deny inter user traffic	Select this checkbox to deny traffic between the clients using this virtual AP profile. The global firewall shown the **Configuration>Advanced Services > Stateful Firewall > Global** window also includes an option to deny all inter-user traffic, regardless of the Virtual AP profile used by those clients. If the global setting to deny inter-user traffic is enabled, all inter-user traffic between clients will be denied, regardless of the settings configured in the virtual AP profiles. If the setting to deny inter-user traffic is disabled globally but enabled on an individual virtual ap, only the traffic between un-trusted users and the clients on that particular virtual AP will be blocked.
Deny time range	Click the drop-down list and select a configured time range for which the AP will deny access. If you have not yet configured a time range, navigate to **Configuration > Security > Access Control > Time Ranges** to define a time range before configuring this setting in the **v**irtual AP profile.
DoS Prevention	If enabled, APs ignore deauthentication frames from clients. This prevents a successful deauthorization attack from being carried out against the AP. This does not affect third-party APs. Default: Disabled
HA Discovery on-association	If enabled, home agent discovery is triggered on client association instead of home agent discovery based on traffic from client. Mobility on association can speed up roaming and improve connectivity for clients that do not send many uplink packets to trigger mobility (VoIP clients). Best practices is to disable this parameter as it increases IP mobility control traffic between controllers in the same mobility domain. Enable this parameter only when voice issues are observed in VoIP clients. Default: Disabled **NOTE:** `ha-disc-onassoc` parameter works only when IP mobility is enabled and configured on the controller. For more information about this parameter, see HA Discovery on Association on page 605
Mobile IP	Enables or disables IP mobility for this virtual AP. Default: Enabled
Preserve Client VLAN	If you select this checkbox, clients retain their previous VLAN assignment if the client disassociates from an AP and then immediately re-outassociates either with same AP or another AP on the same controller.
Remote-AP Operation	Configures when the virtual AP operates on a remote AP: • **always**—Permanently enables the virtual AP (Bridge Mode only). No authentication supported. • **backup**—Enables the virtual AP if the remote AP cannot connect to the controller (Bridge Mode only). No authentication supported. • **persistent**—Permanently enables the virtual AP after the remote AP initially connects to the controller (Bridge Mode only). • **standard**—Enables the virtual AP when the remote AP connects to the controller. Use standard option for tunneled, split-tunneled, and Bridge SSIDs. **NOTE:** Only open/PSK security mode is allowed for always/backup RAP operation. No authentication is supported for always/backup.

Parameter	Description
Station Blacklisting	Select the **Station Blacklisting** checkbox to enable detection of denial of service (DoS) attacks, such as ping or SYN floods, that are not spoofed deauthorization attacks. Default: Enabled
Strict Compliance	If enabled, the AP denies client association requests if the AP and client station have no common rates defined. Some legacy client stations which are not fully 802.11-compliant may not include their configured rates in their association requests. Such non-compliant stations may have difficulty associating with APs unless strict compliance is disabled. This parameter is disabled by default.
VLAN Mobility	Enable or disable VLAN (Layer-2) mobility. Default: Disabled
FDB Update on Assoc	This parameter enables seamless failover for silent clients, allowing them to re-associate. If you select this option, the controller will generate a Layer 2 update on behalf of client to update forwarding tables in bridge devices. Default: Disabled

6. Click **Apply**.

Selective Multicast Stream

The selective multicast group is based only on the packets learned through Internet Group Management Protocol (IGMP).

- When **broadcast-filter all** parameter is enabled, the controller would allow multicast packets to be forwarded only if the following conditions are met:
  - packets originating from the wired side have a destination address range of 225.0.0.0 - 239.255.255.255
  - a station has subscribed to a multicast group.
- When IGMP snooping/proxy is disabled, the controller is not aware of the IGMP membership and drops the multicast flow.
- If DMO is enabled, the packets are sent with 802.11 unicast header.
- If AirGroup is enabled, mDNS (SSDP) packets are sent to the AirGroup application. The common address for mDNS is 224.0.0.251 and SSDP is 239.255.255.250.

Associating Other Profiles to the Virtual AP

Each Virtual AP profile can be associated with the following profile types.

- AAA
- 802.11K
- Handover Trigger Feature Settings
- RRM IE Settings
- Beacon Report Request Settings
- TSM Report Request Settings
- Hotspot 2.0
- SSID
- EDCA Parameters Station
- EDCA Parameters AP
- High-throughput SSID

- 802.11r
- WMM Traffic Management

As a part of the virtual AP profile configuration procedure, you must identify which instance of each profile type associates with the Virtual AP profile. By default, each Virtual AP profile is associated with the **default** versions of the AAA, 802.11k, Hotspot 2.0 and SSID profiles. The Virtual AP profile can also associate with a WMM traffic management profile, but as no WMM profile is associated by default, one must be manually configured.

To configure Virtual AP profile associations:

1. Navigate to **Configuration > Advanced Services > All Profiles**.
2. Expand the **Wireless LAN** menu.
3. Expand the **Virtual AP** menu.
4. Select the Virtual AP you want to configure. The list of associated profile types appears in the **Profiles** list.
5. If a plus [+] sign appears beside an associated profile category, there is more than one profile type in that category. Select that profile category to display the associated profiles within that category.
6. To associate a different profile with the Virtual AP profile, click the name of the any currently associated profile in the **Profiles** list.
7. Click the drop-down list at the top of the **Profile Details** pane and select a different profile to associate to the Virtual AP.
8. Click **Apply**.

Configuring a Virtual AP in the CLI

The example below follows the suggested order of steps to configure a virtual AP using the command-line interface.

```
(host)(config) #vlan 60
!
(host)(config) #ip access-list session THR-POLICY-NAME-WPA2
   user any any permit
!
(host)(config) #user-role THR-ROLE-NAME-WPA2
   session-acl THR-POLICY-NAME-WPA2
!
(host)(config) #aaa authentication dot1x "THR-DOT1X-AUTH-PROFILE-WPA2"
   termination enable
!
(host)(config) #aaa server-group "THR-DOT1X-SERVER-GROUP-WPA2"
   auth-server Internal
!
(host)(config) #aaa profile "THR-AAA-PROFILE-WPA2"
   authentication-dot1x "THR-DOT1X-AUTH-PROFILE-WPA2"
   dot1x-default-role "THR-ROLE-NAME-WPA2"
   dot1x-server-group "THR-DOT1X-SERVER-GROUP-WPA2"
!
(host)(config) #wlan ssid-profile "THR-SSID-PROFILE-WPA2"
   essid "THR-WPA2"
   opmode wpa2-aes
!
(host)(config) #wlan virtual-ap "THR-VIRTUAL-AP-PROFILE-WPA2"
   ssid-profile "THR-SSID-PROFILE-WPA2"
   aaa-profile "THR-AAA-PROFILE-WPA2"
   vlan 60
!
(host)(config) #ap system-profile "THR-AP-SYSTEM-PROFILE"
   lms-ip 1.1.1.1
```

```
   bkup-lms-ip 2.2.2.2
!
(host)(config) #ap-group "THRHQ1-STANDARD"
   virtual-ap "THR-VIRTUAL-AP-PROFILE-WPA2"
(host)(config) #ap-system-profile "THR-AP-SYSTEM-PROFILE"
```

Associating a Virtual AP Profile to an AP or AP Group

Use the following procedures to associate a virtual AP profile to an AP or group of APs.

In the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. Do one of the following:
   - To associate the Virtual AP profile to a single AP, click the **AP specific** tab, and select the AP.
   - To associate the Virtual AP profile to an AP group, click the **AP Group** tab, and select the AP Group.
3. In the **Profiles** list, expand the **Wireless LAN** menu.
4. Select **Virtual AP**. The Profile Details window dislays the Virtual AP profiles currently associated to the AP or AP group.
5. Click the **Add a Profile** drop-down list and select a new Virtual AP profile to associate to the AP. You can associate multiple AP profiiles to an AP, but each virtual AP profile must reference a SSID profile with a different network name (SSID).
6. Click **Apply**



Although you can create mutliple Virtual AP profiles that reference a single SSID profile, only one of these profiles can be applied to an AP or AP group.

In the CLI

```
(host) (config) ap-group <ap-group>
virtual-ap <vap-profile>


(host) (config) ap-name <ap-name>
virtual-ap <vap-profile>
```

Excluding a Virtual AP Profile

You can exclude one or more virtual AP profiles from an individual AP. This prevents a virtual AP, defined at the AP group level, from being applied to a specific AP. For example, you can apply the virtual AP profile that corresponds to the “Corpnet” SSID to the “default” AP group. If you do not want the “Corpnet” SSID to be advertised on the AP in the lobby, you can specify the virtual AP profile that contains the “Corpnet” SSID configuration be excluded from that AP.

In the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration > AP Specific** page.
2. Do one of the following:
   - If the AP you want to exclude is included in the list, click **Edit** for the AP.
   - If the AP does not appear in the list, click **New.** Either type in the name of the AP, or select the AP from the drop-down list. Then click **Add**.
3. Select **Wireless LAN** under the Profiles list, then select **Excluded Virtual AP.**

4. Select the name of the virtual AP profile you want to exclude from the drop down menu (under **Profile Details**) and click **Add.** The profile name appears in the **Excluded Virtual APs** list. You can add multiple profile names in the same way.
5. To remove a profile name from the Excluded Virtual APs list, select the profile name and click **Delete.**
6. Click **Apply.**

In the CLI

```
(host)(config) #ap-name <name>
   exclude-virtual-ap <profile>
```

Virtual AP Configuration Workflow

The following workflow lists the tasks to configure a virtual AP that uses 802.1X authentication. Click any of the links below for details on the configuration procedures for that task.

Using the WebUI

1. Configure your authentication servers.
2. Create an authentication server group, and assign the authentication servers you configured in step 1 to that server group.
3. Configure a firewall access policy for a group of users
4. Create a user role, and assign the firewall access policy you created in step 3 to that user role.
5. Create an AAA profile.
   a. Assign the user role defined in step 4 to the AAA profile's **802.1X Authentication Default Role**
   b. Associate the server group you created in step 2 to the AAA profile.
6. Create a new SSID profile
7. Create a new virtual AP profile.
8. Associate the virtual AP profile to the AAA profile you created in Step 5.
9. Associate the virtual AP profile to the SSID profile you created in Step 6.

Using the CLI

The example below follows the suggested order of steps to configure a virtual AP using the command-line interface.

```
(host)(config) #aaa server-group "THR-DOT1X-SERVER-GROUP-WPA2"
   auth-server Internal
!
ip access-list session THR-POLICY-NAME-WPA2
   user any any permit
!
(host)(config) #user-role THR-ROLE-NAME-WPA2
   session-acl THR-POLICY-NAME-WPA2
!
(host)(config) #aaa server-group "THR-DOT1X-SERVER-GROUP-WPA2"
   auth-server Internal
!
(host)(config) #aaa profile "THR-AAA-PROFILE-WPA2"
   dot1x-default-role "THR-ROLE-NAME-WPA2"
   dot1x-server-group "THR-DOT1X-SERVER-GROUP-WPA2"
!
(host)(config) #wlan ssid-profile "THR-SSID-PROFILE-WPA2"
   essid "THR-WPA2"
   opmode wpa2-aes
```

```
!
(host)(config) #wlan virtual-ap "THR-VIRTUAL-AP-PROFILE-WPA2"
   ssid-profile "THR-SSID-PROFILE-WPA2"
   aaa-profile "THR-AAA-PROFILE-WPA2"
   vlan 60
!
(host)(config) #ap-group "THRHQ1-STANDARD"
   virtual-ap "THR-VIRTUAL-AP-PROFILE-WPA2"
```

Radio Resource Management (802.11k)

The 802.11k protocol provides mechanisms for APs and clients to dynamically measure the available radio resources. In an 802.11k enabled network, APs and clients can send neighbor reports, beacon reports, and link measurement reports to each other. This allows the APs and clients to take appropriate connection actions.



The handover process is available for voice clients that support the 802.11k standard and have the ability to transmit and receive beacon reports. For information on configuring the handoff trigger feature, see Enabling Wi-Fi Edge Detection and Handover for Voice Clients on page 922

This topic includes the following procedures:

- Configuring the 802.11k Profile
- Configuring Radio Resource Management Information Elements
- Configuring Beacon Report Requests
- Configuring Traffic Stream Measurement Report Requests

Configuring the 802.11k Profile

The following procedures outline the steps to configure 802.11k parameters.

In the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab.
   - If you selected the **AP Group** tab, click the **Edit** button by the AP group name for which you want to configure the new 802.11k profile.
   - If you selected the **AP Specific** tab, click the **Edit** button by the AP for which you want to create the 802.11K profile.
2. In the Profiles list, expand the **Wireless LAN** menu, then expand the **Virtual AP** menu.
3. Select the Virtual AP profile for which you want to configure 802.11k settings.

   To edit an existing 802.11k profile, click the **802.11K Profile** drop-down list In the **Profile Details** window pane and select the 802.1x profile you want to edit.

   or

   To create a new 802.11k profile, click the **802.11k Profile** drop-down list and select **New**. Enter a new 802.11k profile name in the field to the right of the drop-down list.
4. Configure your 802.11k radio settings. Table 73 outlines the parameters you can configure in the 802.11k profile. Click **Apply** to save your settings.

Table 73: *802.11k Profile Parameters*

Parameter	Description
Advertise 802.11k Capability	Select this option to allow Virtual APs using this profile to advertise 802.11k capability. Default: Disabled
Forcefully disassociate on-hook voice clients	Select this option to allow the AP to forcefully disassociate *on-hook* voice clients (clients that are not on a call) after period of inactivity. Without the forced disassociation feature, if an AP has reached its call admission control limits and an on-hook voice client wants to start a new call, that client may be denied. If forced disassociation is enabled, those clients can associate to a neighboring AP that can fulfill their QoS requirements. Default: Disabled
Measurement Mode for Beacon Reports	Click the **Measurement Mode for Beacon Reports** drop-down list and specify one of the following measurement modes: • **active**—Enables active beacon measurement mode. In this mode, the client sends a probe request to the broadcast destination address on all supported channels, sets a measurement duration timer, and, at the end of the measurement duration, compiles all received beacons or probe response with the requested SSID and BSSID into a measurement report. • **beacon-table**—Enables beacon-table beacon measurement mode. In this mode, the client measures beacons and returns a report with stored beacon information for any supported channel with the requested SSID and BSSID. The client does not perform any additional measurements. • **passive**—Enables passive beacon measurement mode. In this mode, the client sets a measurement duration timer, and, at the end of the measurement duration, compiles all received beacons or probe response with the requested SSID and BSSID into a measurement report. **NOTE:** If a station doesn't support the selected measurement mode, it returns a Beacon Measurement Report with the Incapable bit set in the Measurement Report Mode field. Default Mode: beacon-table
Channel for Beacon Requests in 'A' band	This value is sent in the 'Channel' field of the beacon requests on the 'A' radio. You can specify values in the range 34 to 165. The default value is 36.
Channel for Beacon Requests in 'BG' band	This value is sent in the 'Channel' field of the Beacon Requests on the 'BG' radio. You can specify values in the range 1 to 14. The default value is 1.
Channel for AP Channel Reports in 'A' band	This value is sent in the 'Channel' field of the AP channel reports on the 'A' radio. You can specify values in the range 34 to 165. The default value is 36.
Channel for AP Channel Reports in 'BG' band	This value is sent in the 'Channel' field of the AP channel reports on the 'BG' radio. You can specify values in the range 1 to 14. The default value is 1.
Time duration between consecutive Beacon Requests	This option configures the time duration between two consecutive beacon requests sent to a dot11K client. By default, the beacon requests are sent to a dot11K client every 60 seconds. However, if a different value is required, the `bcn-req-time` option can be used.

Parameter	Description
	This permits values in the range from 10 seconds to 200 seconds. A value of 0 is used to indicate that the generation of Beacon Request frames is turned off.
Time duration between consecutive Link Measurement Requests	This option configures the time duration between two consecutive link measurement requests sent to an dot11K client. By default, link measurement requests are sent to a dot11K client every 61 seconds. This parameter permits values in the range from 10 seconds to 200 seconds. A value of 0 is used to indicate that the generation of Link Measurement Request frames is turned off.
Time duration between consecutive Transmit Stream Measurement Request	This option configures the time duration between two consecutive transmit stream measurement requests sent to a dot11K client. By default, the transmit stream measurement requests are sent to a dot11K client every 90 seconds. This permits values in the range from 10 seconds to 200 seconds. A value of 0 is used to indicate that the generation of Transmit Stream Measurement Request frames is turned off.
Handover Trigger Feature Settings Profile	This command configures a Handover Trigger Profile. This profile consists of the configurable parameters for the 'Wi-Fi Edge Detection and Handover of Voice Clients' feature.
Beacon Report Request Settings Profile	Configure a Beacon Report Request Profile to provide the parameters for the Beacon Report Request frames.
TSM Report Request Settings Profile	This command configures a TSM Report Request Profile which is used to provide values to the Transmit Stream/Category Measurement Request frame.

In the CLI

Use the following command to configure 802.11k profiles. The available parameters for this profile are described in Table 73.

```
wlan dotllk <profile-name>
```

Configuring Radio Resource Management Information Elements

ArubaOS supports the following radio resource management information elements (RRM IEs) for APs with 802.11k support enabled. These settings can be enabled through the WebUI or CLI.

In the WebUI

To select the RRM IEs to be sent in beacons and probe responses using the WebUI:

1. Navigate to **Configuration>Advanced Services>All Profile Management**.
2. Expand the **Wireless LAN** menu and select **RRM IE**.
3. Select the RRM IE profile you want to configure, then select any of the following IE types to enable that information element in beacons and probe responses. (All IE types are sent by default.)

Table 74: *RRM IE Parameters*

Parameter	Description
Advertise Enabled Capabilities IE	This value is used to determine if the RRM Enabled Capabilities IE should be advertised in the beacon frames. A value of "Enabled" allows the RRM Enabled Capabilities IE to be present in the beacon frames when 802.11K capability is enabled. A value of "Disabled" prevents the advertisement of the RRM Enabled Capabilities IE in the beacon frames when 802.11K capability is enabled.
Advertise Country IE	This value is used to determine if the Country IE should be advertised in the beacon frames. A value of "Enabled" allows the Country IE to be present in the beacon frames when 802.11K capability is enabled. A value of "Disabled" prevents the advertisement of the Country IE in the beacon frames when 802.11K capability is enabled.
Advertise Power Constraint IE	This value is used to determine if the Power Constraint IE should be advertised in the beacon frames. A value of "Enabled" allows the Power Constraint IE to be present in the beacon frames when 802.11K capability is enabled. A value of "Disabled" prevents the advertisement of the Power Constraint IE in the beacon frames when 802.11K capability is enabled.
Advertise TPC Report IE	This value is used to determine if the TPC Report IE should be advertised in the beacon frames. A value of "Enabled" allows the TPC Report IE to be present in the beacon frames when 802.11K capability is enabled. A value of "Disabled" prevents the advertisement of the TPC Report IE in the beacon frames when 802.11K capability is enabled.
Advertise QBSS Load IE	This value is used to determine if the QBSS Load IE should be advertised in the beacon frames. A value of "Enabled" allows the QBSS Load IE to be present in the beacon frames when 802.11K capability is enabled. A value of "Disabled" prevents the advertisement of the QBSS Load IE in the beacon frames when 802.11K capability is enabled. The default value is "Enabled".
Advertise BSS AAC IE	This value is used to determine if the BSS Available Admission Capacity IE should be advertised in the beacon frames. A value of "Enabled" allows the BSS Available Admission Capacity IE to be present in the beacon frames when 802.11K capability is enabled. A value of "Disabled" prevents the advertisement of the BSS Available Admission Capacity IE in the beacon frames when 802.11K capability is enabled.
Advertise Quiet IE	This value is used to determine if the Quiet IE should be advertised in the beacon frames. A value of "Enabled" allows the Quiet IE to be present in the beacon frames when 802.11K capability is enabled. A value of "Disabled" prevents the advertisement of the Quiet IE in the beacon frames when 802.11K capability is enabled.

4. Click **Apply Changes** to save your settings.

In the CLI

To use the CLI to configure radio resource management information elements in the RRM IE profile, access the CLI in config mode and issue the following commands:

```
wlan rrm-ie-profile <profile>
   bss-aac-ie
   clone
   country-ie
   enabled-capabilities-ie
   no...
   pwr-constraint-ie
```

```
qbss-load-ie
quiet-ie
tpc-report-ie
```

Configuring Beacon Report Requests

The beacon report requests are sent only to 802.11k-compliant clients that advertise Beacon Report Capability in their RRM Enabled Capabilities IE. The beacon request frames are sent every 60 seconds.

The content of the report requests can be defined in the Beacon Report Request profile using the WebUI or CLI.

In the WebUI

To select the information to be sent in beacon report reqeusts using the WebUI:

1. Navigate to **Configuration>Advanced Services>All Profile Management**.
2. Expand the **Wireless LAN** menu and select **Beacon Report Request**.
3. Select the Beacon Report Request profile you want to configure.
4. Define the settings described in the table below, then click **Apply Changes** to save your settings.

Table 75: *Beacon Report Request Settings*

Parameter	Description
Interface	This field is used to specify the Radio interface for transmitting the Beacon Report Request frame. It can have a value of either 0 or 1. The default value is 1.
Regulatory Class	This option is used to specify the Regulatory Class field in the Beacon Report Request frame. It can be set to one of the following: - • **5** (for 5 GHz band) • **12** (for 2.4 GHz band)
Channel	This option is used to set the Channel field in the Beacon Report Request frame. The Channel value can be set to one of the following: - the channel of the AP (when Measurement Mode is set to either 'Passive' or 'Active-All channels') - 0 (when Measurement Mode is set to 'Beacon Table') - 255 (when Measurement Mode is set to 'Active-Channel Report')
Randomization Interval	This value is used to set the Randomization Interval field in the Beacon Report Request frame. The Randomization Interval is used to specify the desired maximum random delay in the measurement start time. It is expressed in units of TUs (Time Units). A Randomization Interval of 0 in a measurement request indicates that no random delay is to be used. This field can be given a value in the range (0, 65535). The default value is 0.
Measurement Duration	This value is used to set the Measurement Duration field in the Beacon Report Request frame. The Measurement Duration is set to the duration of the requested measurement. It is expressed in units of TUs. This field can be given a value in the range (0, 65535). The default value is 0.
Measurement Mode for Beacon Reports	Click the **Measurement Mode for Beacon Reports** drop-down list and specify one of the following measurement modes: • **active**—Enables active beacon measurement mode. In this mode, the client

Parameter	Description
	sends a probe request to the broadcast destination address on all supported channels, sets a measurement duration timer, and, at the end of the measurement duration, compiles all received beacons or probe response with the requested SSID and BSSID into a measurement report. • **beacon-table**—Enables beacon-table beacon measurement mode. In this mode, the client measures beacons and returns a report with stored beacon information for any supported channel with the requested SSID and BSSID. The client does not perform any additional measurements. • **passive**—Enables passive beacon measurement mode. In this mode, the client sets a measurement duration timer, and, at the end of the measurement duration, compiles all received beacons or probe response with the requested SSID and BSSID into a measurement report. **NOTE:** If a station doesn't support the selected measurement mode, it returns a Beacon Measurement Report with the Incapable bit set in the Measurement Report Mode field. Default Mode: beacon-table
Reporting Condition	This option is used to indicate the value for the "Reporting Condition" field in the Beacon Reporting Information sub-element present in the Beacon Report Request frame. It can have a range from 0 to 255. The default value is 0.
ESSID name	This option is used to indicate the value for the "SSID" field in the Beacon Report Request frame. It corresponds to the SSID Name for which the Beacon Report Request frame needs to be generated. It is a string with a minimum length of 1 and a maximum length of 32.
Reporting Detail	This option is used to indicate the value for the "Detail" field in the Reporting Detail sub-element present in the Beacon Report Request frame. It is set to "Disabled" by default.
Measurement Duration Mandatory	This value is used to set the "Duration Mandatory" bit of the Measurement Request Mode field of the Beacon Report Request frame. The default value is "Disabled".
Request Information values	This option is used to indicate the contents of the Request Information IE that could be present in the Beacon Report Request frame. The Request Information IE is present for all Measurement Modes except the 'Beacon Table' mode. It consists of a list of Element IDs that should be included by the client in the response frame.

In the CLI

To select the information to be sent in beacon report requests using the command-line interface, access the CLI in config mode and issue the following commands.

```
wlan bcn-rpt-req-profile <profile>
```

Configuring Traffic Stream Measurement Report Requests

The Traffic Stream Measurement(TSM) report requests are sent only to dot11k compliant clients that advertise a traffic stream report capability. The TSM report request frames are sent every 60 seconds. The content of the report requests can be defined in the TSM Report Request profile using the WebUI or CLI.

In the WebUI

To select the information to be sent in TSM report reqeusts using the WebUI:

1. Navigate to **Configuration > Advanced Services > All Profile Management**.
2. Expand the **Wireless LAN** menu and select **TSM Report Request**.
3. Select the TSMReport Request profile you want to configure.
4. Define the settings described in the table below, then click **Apply Changes** to save your settings.

Table 76: *TSM Report Request Settings*

Parameter	Description
Request Mode for TSM Report Request	Select one of the following request modes: • normal • triggered This value is used to determine the request mode for the Transmit Stream/Category Measurement Request frame. A Transmit Stream/Category Measurement Request frame can be sent in either normal mode or triggered mode. There are two options for this parameter **normal** and **triggered**. When the **triggered** option is selected, the Transmit Stream/Category Measurement Request frame is sent only when the trigger condition occurs. The default value for this field is **normal**.
Number of repetitions	This value is used to set the "Number of Repetitions" field in the Transmit Stream/Category Measurement Request frame. The **Number of Repetitions** field contains the requested number of repetitions for all the Measurement Request elements in this frame. A value of zero in this field indicates Measurement Request elements are executed once without repetition. A value of 65535 in the Number of Repetitions field indicates Measurement Request elements are repeated until the measurement is cancelled or superseded. This field has values in the range (0, 65535). The default value is 65535.
Duration Mandatory	This value is used to set the "Duration Mandatory" bit of the Measurement Request Mode field of the Transmit Stream/Category Measurement Request frame. The default value is **enabled**.
Randomization Interval	This value is used to set the Randomization Interval field in the Transmit Stream/Category Measurement Request frame. The Randomization Interval is used to specify the desired maximum random delay in the measurement start time. It is expressed in units of TUs (Time Units). When the request mode for the Transmit Stream/Category Measurement Request frame is set to "triggered", the Randomization Interval is not used and is set to 0. A Randomization Interval of 0 in a measurement request indicates that no random delay is to be used. This field can be given a value in the range (0, 65535). The default value is 0.
Measurement Duration	This value is used to set the Measurement Duration field in the

Parameter	Description
	Transmit Stream/Category Measurement Request frame. The Measurement Duration is set to the duration of the requested measurement. It is expressed in units of TUs. When the request mode for the Transmit Stream/Category Measurement Request frame is set to **triggered**, the Measurement Duration field should be set to 0. This field can be given a value in the range (0, 65535). The default value is 9776.
Traffic ID	The value is used to set the Traffic Identifier field in the Transmit Stream/Category Measurement Request frame. The Traffic Identifier field contains the TID subfield. The TID subfield indicates the TC or TS for which traffic is to be measured. This field can be given a value in the range (0, 255). The default value is 96
Bin 0 Range	This value is used to set the 'Bin 0 Range' field in the Transmit Stream/Category Measurement Request frame. Bin 0 Range indicates the delay range of the first bin (Bin 0) of the Transmit Delay Histogram, expressed in units of TUs. This field can be given a value in the range (0, 255). The default value is 6.

In the CLI

To select the information to be sent in TSM report requests using the command-line interface, access the CLI in config mode and issue the following commands.

```
wlan tsm-req-profile <default>
```

BSS Transition Management (802.11v)

BSS Transition Management enables an AP to request a voice client to transition to a specific AP, or suggest a set of preferred APs to a voice client, due to network load balancing or BSS termination. This helps the voice client identify the best AP to which that client should transition to as that client roams. ArubaOS supports BSS Transition Management features defined by the 802.11v standard.

The BSS Transition capability can improve throughput, data rates and QoS for the voice clients in a network by shifting (via transition) the individual voice traffic loads to more appropriate points of association within the ESS.

Frame Types

BSS Transition Management uses the following frame types:

- **Query:** A Query frame is sent by the voice client that supports BSS transition management requesting a BSS transition candidate list to its associated AP, if the associated AP indicates that it supports the BSS transition capability.
- **Request:** An AP that supports BSS Transition Management responds to a BSS Transition Management Query frame with a BSS Transition Management Request frame. The AP may also send an unsolicited BSS Transition Management Request frame to a voice client at any time, if the client supports the BSS Transition Management capability. The Request frame also contains a Disassociation flag. If the flag is set, then the AP forcefully disassociates the client after 10 beacon intervals.

- **Response:** A Response frame is sent by the voice client back to the AP, informing whether it accepts or denies the transition.

802.11k and 802.11v clients

For 802.11k capable clients, the client management framework uses the actual beacon report generated by the client in response to a beacon report request sent by the AP. This beacon report replaces the virtual beacon report for that client.

For 802.11v capable clients, the controller uses the 802.11v BSS Transition message to steer clients to the desired AP upon receiving a client steer trigger from the AP.

Enabling Vss

802.11To enable 802.11v BSS transition management, enable the **Advertise 802.11k Capability** parameter in an 802.11k profile, then ensure that 802.11k profile is associated to a Virtual AP profile.

For more information on the 802.11k profile, see Radio Resource Management (802.11k) on page 404.

Fast BSS Transition (802.11r)

ArubaOS provides support for Fast BSS Transition as part of the 802.11r implementation. Fast BSS Transition mechanism minimizes the delay when a voice client transitions from one BSS to another within the same ESS. Fast BSS Transition establishes security and QoS states at the target AP before or during a re-association. This minimizes the time required to resume data connectivity when a BSS transition happens.

The following table provides the modes in which Fast BSS Transition is supported:

Table 77: *Supported VAP Forwarding Modes*

VAP Forwarding Mode	Support for 802.11r
Tunnel Mode	Yes
Decrypt-Tunnel Mode	Yes
Split-Tunnel Mode	No
Bridge Mode	Beta quality

Important Points to Remember

- Fast BSS Transition is operational only if the wireless client has support for 802.11r standard. If the client does not have support for 802.11r standard, it falls back to normal WPA2 authentication method.
- If dot11r is enabled, iOS clients such as iPad/iPhone gen1 (limitation on iOS) and all MAC-OS clients (limitation on MAC) fail to connect to the network.

Configuring Fast BSS Transition

You can enable and configure Fast BSS Transition on a per Virtual AP basis. You must create an 802.11r profile and associate that with the Virtual AP profile through an SSID profile. You can create and configure an 802.11r profile using the WebUI or CLI.



Fast BSS transition is operational only with WPA2-Enterprise or WPA2-Personal.

In the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration**window. Select either the **AP Group** or **AP Specific** tab.
   a. If you selected the **AP Group** tab, click the AP group name for which you want to configure the 802.11R profile.
   b. If you selected the **AP Specific** tab, click the AP for which you want to configure the 802.11R profile.
2. In the **Profiles** list, expand the **Wireless LAN**menu, then expand the **Virtual AP**menu.
3. Select the Virtual AP profile for which you want to configure the 802.11r settings and expand **SSID Profile**.
4. Select the SSID profile on which you want to configure the 802.11r settings and select **802.11R Profile**.
   a. To edit an existing 802.11r profile, click the **802.11R Profile**drop-down list in the **Profile Details**window pane and select the 802.11r profile you want to edit.

   or

   b. To create a new 802.11r Profile, click the **802.11R Profile** drop-down list and select **New**. Enter a new 802.11r profile name in the field to the right of the drop-down list.



You cannot use spaces in profile names.

5. Configure the following 802.11r radio settings.
   a. Select the **Advertise 802.11r Capability** option to allow Virtual APs using this profile to advertise 802.11r capability.
   b. Enter the mobility domain ID value (1-65535) in the **802.11r Mobility Domain ID** field. The default value is 1.
   c. Enter the R1 Key timeout value in seconds (60-86400) for decrypt-tunnel or bridge mode in the **802.11r R1 Key Duration** field. The default value is 3600.
6. Click**Apply** to save your settings.

In the CLI

Create an 802.11r profile using the following command:

```
(host) (config) #wlan dot11r-profile voice-enterprise
```

Enable Fast BSS Transition using the following command:

```
(host) (802.11R Profile "voice-enterprise") #dot11r
```

Configure a mobility domain ID that uniquely identifies a mobility domain using the following command:

```
(host) (802.11R Profile "voice-enterprise") #mob-domain-id <1-65535>
```

The default value is 1.

Configure the r1 key timeout value in seconds for decrypt-tunnel or bridge mode using the following command:

```
(host) (802.11R Profile "voice-enterprise") #key_duration <60-86400>
```

The default value is 3600 seconds.

Apply the 802.11r profile to an SSID profile using the following command:

```
(host) (config) #wlan ssid-profile voice dot11r-profile voice-enterprise
```

You can advertise the 802.11r capability on the Virtual AP profile by applying the SSID profile. Use the following command to apply the SSID profile to the Virtual AP profile:

```
(host) (config) #wlan virtual-ap voice-AP ssid-profile voice
```

Troubleshooting Fast BSS Transition

ArubaOS provides various troubleshooting options to verify the Fast BSS Transition functionalities.

In decrypt-tunnel mode and bridge mode, each r0 key generates up to four r1 keys and the controller pushes each r1 key to the corresponding AP. A few commands are added to help verifying the pushing functionality:

Execute the following command to view all the r1 keys that are stored in an AP:

```
(host)(config) #show ap debug dot11r state

   [ap-name <ap-name> | ip-addr <ip-addr>]
```

You can filter the output based on the AP name, BSSID, or IP address.

```
(host)(config) #show ap debug dot11r state ap-name MAcage-105-GL
   Stored R1 Keys
   --------------
   Station MAC        Mobility Domain ID  Validity Duration  R1 Key
   -----------        ------------------  -----------------  ------
   00:50:43:21:01:b8  1                   3568               (32): 94 ff 18 0a 5f 47 8b 3e 95
   2b 93 31 bd
    44 58 fe fe 6a ad aa 1d d7 29 94 fb 5b 7c 15 76 66 d2 1f
```

You can use the following command to remove an r1 key from an AP when the AP does not have a cached r1 key during Fast BSS Transition roaming.

```
(host) #ap debug dot11r remove-key <sta-mac> ap-name <ap-name> | ip-addr <ip-addr>


(host) #ap debug dot11r remove-key 00:50:43:21:01:b8 ap-name MAcage-105-GL
```

Execute the following command to check if the r1 key is removed from the AP:

```
(host)(config) #show ap debug dot11r state ap-name MAcage-105-GL


Stored R1 Keys
--------------
Station MAC  Mobility Domain ID  Validity Duration  R1 Key
-----------  ------------------  -----------------  ------
```

Execute the following command to view the hit/miss rate of r1 keys cached on an AP before a Fast BSS Transition roaming. This counter helps to verify if enough r1 keys are pushed to the neighboring APs.

```
(host)(config) #show ap debug dot11r efficiency <client-mac>
(host)(config) #show ap debug dot11r efficiency


Fast Roaming R1 Key Efficiency
------------------------------
Client MAC         Hit (%)  Miss (%)
----------         -------  --------
00:50:43:21:01:b8  0 (0%)   0 (0%)
```

SSID Profiles

A Service Set Identifier (SSID) is the network or WLAN that any client sees. A SSID profile defines the name of the network, authentication type for the network, basic rates, transmit rates, SSID cloaking, and certain WMM settings for the network.

SSID Profile Overview

ArubaOS supports different types of the Advanced Encryption Standard (AES), Temporal Key Integrity Protocol (TKIP), and wired equivalent privacy (WEP) encryption. AES is the most secure and recommended encryption method. Most modern devices are AES capable and AES should be the default encryption method. Use TKIP only when the network includes devices that do not support AES. In these situations, use a separate SSID for devices that are only capable of TKIP.

Suite-B Cryptography

The Suite-B (bSec) protocol is a pre-standard protocol that has been proposed to the IEEE 802.11 committee as an alternative to 802.11i. The main difference between bSec and standard 802.11i is that bSec implements Suite-B algorithms wherever possible. Notably, AES-CCM is replaced by AES-GCM, and the Key Derivation Function (KDF) of 802.11i is upgraded to support SHA-256 and SHA-384. In order to provide interoperability with standard Wi-Fi software drivers, bSec is implemented as a shim layer between standard 802.11 Wi-Fi and a Layer 3 protocol such as IP. A controller configured to advertise a bSec SSID will advertise an open network, however only bSec frames will be permitted on the network.



This feature requires the ACR license.

The bSec protocol requires that you use VIA 2.1.1 or greater on the client device. Consult VIA documentation for more information on configuring and installing VIA.

The bSec protocol is available in 128-bit mode and 256-bit mode. The number of bits specifies the length of the AES-GCM encryption key. Using United States Department of Defense classification terminology, bSec-128 is suitable for protection of information up to the SECRET level, while bSec-256 is suitable for protection of information up to the TOP SECRET level.

Suite-B AES-128-GCM and AES-256-GCM encryption is supported by the ArubaOS hardware. Note, however, that not all controllers support Suite-B encryption. The table below describes the controller support for Suite-B encryption in ArubaOS.

Controller	Serial Number Prefix	ACR License Support
W-7200 Series	All serial numbers supported	Yes
W-600 Series	All serial numbers supported	Yes
W-3000 Series	FC	Yes
W-3000 Series	F	No
W-6000M3 card	AK	Yes
W-6000M3 card	A	No

To determine the serial number prefix for your controller, issue the CLI command **show inventory** and note the prefix before the system serial number. The serial number prefix in the example below appears in **bold.**

```
(host) #show inventory
Supervisor Card slot            : 0
System Serial#                  : AK0093676
SC      Assembly#               : 2010052B (Rev:02.01)
SC      Serial#                 : F01629529 (Date:03/29/10)
SC      Model#                  : W-3600-US
```

Wi-Fi Multimedia Protection

Wi-Fi Multimedia (WMM®) is a Wi-Fi Alliance® certification program that is based on the IEEE 802.11e amendment. WMM ensures QoS for latency-sensitive traffic in the air. WMM divides the traffic into four queues or access categories:

- voice
- video
- best effort
- background

Management Frame Protection

ArubaOS supports the IEEE 802.11w standard, also known as Management Frame Protection (MFP). MFP makes it difficult for an attacker to deny service by spoofing Deauth and Disassoc management frames. MFP uses 802.11i (Robust Security Network) framework that establishes encryption keys between the client and AP.

MFP is configured on a virtual AP (VAP) as part of the wlan ssid-profile. There are two parameters that can be configured, mfp-capable and mfp-required. Both are disabled by default.



MFP can only be enabled on SSIDs that support WPA2. MFP is not supported on virtual APs using tunnel forwarding mode.

Configuring the SSID Profile

Follow the procedures below to create a new SSID profile and associate that profile to your Virtual AP.

In the WebUI

1. Navigate to **Advanced Services > All Profile Management**.
2. In the **Profiles** list, expand the **Wireless LAN** menu, then select **SSID**.
3. Select an existing profile from the Profile Details pane, or enter create a new profile by entering a new name into the entry blank, then clicking **Add**.
4. Configure the SSID profile parameters described in Table 51, then click **Apply**.

The SSID profile configuration settings are divided into two tabs, **Basic** and **Advanced.** The **Basic** tab displays only those configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. If you change a setting on one tab then click and display the other tab without saving your configuration, that setting will revert to its previous value.

Table 78: *SSID Profile Parameters*

Parameter	Description
Basic SSID Profile Settings	
Network Name	Name that uniquely identifies a wireless network. The network name, or *ESSID* can be up to 31 characters. If the ESSID includes spaces, you must enclose it in quotation marks.
Network Authentication	The layer-2 authentication to be used on this ESSID to protect access and ensure the privacy of the data transmitted to and from the network. • None

Parameter	Description
	• 802.1x/WEP • WPA • WPA-PSK • WPA2 • WPA2-PSK • xSec • Mixed If you select the **Mixed** authentication option, a drop-down list will appear in the Network Authentication section. Click this drop-down list and select the combination of authentication types supported by APs using this SSID profile.
Encryption	This field shows the default encryption type used on this ESSID. Unselect the default encryption type if you do not want encryption, or click the **Advanced** tab to define a new encryption type.
Keys	If you selected WPA-PSK or WPA2-PSK authentication or a mixed authentication type that supports pre-shared keys, enter and confirm the Hex Key or PSK passphrase in the **PSK Key/Passphrase** and **Confirm PSK Key/Passphrase** fields. • To define a hex key, enter a 64-character hexadecimal string. • To define a PSK passphrase, enter san ASCII string 8-63 characters in length. Next click the **Format** drop-down list and select **Hex** or **PSK** Passphrase to select the format for the key or passphrase.
Advanced SSID Profile Settings	
SSID Enable	Click this checkbox to enable or disable the SSID. The SSID is enabled by default.
Encryption	Select one of the following encryption types
xSec	Encryption and tunneling of Layer-2 traffic between the controller and wired or wireless clients, or between controllers. To use xSec encryption, you must use a RADIUS authentication server. For clients, you must install the Funk Odyssey client software. Requires installation of the xSec license. For xSec between controllers, you must install an xSec license in each controller.
opensystem	No authentication and encryption.
static-wep	WEP with static keys.
dynamic-wep	WEP with dynamic keys.
wpa-tkip	WPA with TKIP encryption and dynamic keys using 802.1x.
wpa-aes	WPA with AES encryption and dynamic keys using 802.1x.
wpa-psk-tkip	WPA with TKIP encryption using a preshared key.
wpa-psk-aes	WPA with AES encryption using a preshared key.
wpa2-aes	WPA2 with AES encryption and dynamic keys using 802.1x.
wpa2-psk-aes	WPA2 with AES encryption using a preshared key.
wpa2-psk-tkip	WPA2 with TKIP encryption using a preshared key.

Parameter	Description
wpa2-tkip	WPA2 with TKIP encryption and dynamic keys using 802.1x.
wpa2-aes-gcm-128	WPA2 with AES GCM-128 (Suite-b) encryption and dynamic keys using 802.1X. **NOTE:** This parameter requires the ACR license. For further information on Suite-B encryption, see SSID Profiles on page 414.
wpa2-aes-gcm-256	WPA2 with AES GCM-256 (Suite-b) encryption and dynamic keys using 802.1X. **NOTE:** This parameter requires the ACR license. For further information on Suite-B encryption, see SSID Profiles on page 414.
Enable Management Frame Protection	When selected, the SSID supports MFP-capable and traditional clients. **NOTE:** MFP can only be enabled on SSIDs that support WPA2.
Require Management Frame Protection	When selected, the SSID supports MFP-capable clients only. **NOTE:** MFP can only be enabled on SSIDs that support WPA2.
DTIM Interval	Specifies the interval, in milliseconds, between the sending of Delivery Traffic Indication Messages (DTIMs) in the beacon. This is the maximum number of beacon cycles before unacknowledged network broadcasts are flushed. When using wireless clients that employ power management features to sleep, the client must revive at least once during the DTIM period to receive broadcasts
802.11g Transmit Rates	Select the set of 802.11b/g rates at which the AP is allowed to send data. The actual transmit rate depends on what the client is able to handle, based on information sent at the time of association and on the current error/loss rate of the client.
802.11g Basic Rates	Select the set of supported 802.11b/g rates that are advertised in beacon frames and probe responses.
802.11a Transmit Rates	Select the set of 802.11a rates at which the AP is allowed to send data. The actual transmit rate depends on what the client is able to handle, based on information sent at the time of association and on the current error/loss rate of the client.
802.11a Basic Rates	Select the set of supported 802.11a rates, in Mbps, that are advertised in beacon frames and probe responses.
Station Ageout Time	Time, in seconds, that a client is allowed to remain idle before being aged out.
Max Transmit Attempts	Maximum number of retries allowed for the AP to send a frame.
RTS Threshold	Wireless clients transmitting frames larger than this threshold must issue Request to Send (RTS) and wait for the AP to respond with Clear to Send (CTS). This helps prevent mid-air collisions for wireless clients that are not within wireless peer range and cannot detect when other wireless clients are transmitting. The default value is 2333 bytes.
Short Preamble	Click this checkbox to enable or disable a short preamble for 802.11b/g radios. Network performance may be higher when short preamble is enabled. In mixed radio environments, some 802.11b wireless client stations may experience difficulty associating with the AP using short preamble. To use only long preamble, disable short preamble. Legacy client devices that use only long preamble generally can be updated to support short preamble.

Parameter	Description
Max Associations	Maximum number of wireless clients for the AP. The supported range is 0-256 clients.
Wireless Multimedia (WMM)	Enables or disables WMM, also known as IEEE 802.11e Enhanced Distribution Coordination Function (EDCF). WMM provides prioritization of specific traffic relative to other traffic in the network.
Wireless Multimedia U-APSD (WMM-UAPSD) Powersave	Enable Wireless Multimedia (WMM) UAPSD powersave.
WMM TSPEC Min Inactivity Interval	Specify the minimum inactivity time-out threshold of WMM traffic. This setting is useful in environments where low inactivity interval time-outs are advertised, which may cause unwanted timeouts. The supported range is 0-3,600,000 milliseconds, and the default value is 0 milliseconds.
Override DSCP mappings for WMM clients	Override the default DSCP mappings in the SSID profile with the ToS value. This setting is useful when you want to set a non-default ToS value for a specific traffic.
DSCP mapping for WMM voice AC	DSCP used to map WMM voice traffic. The supported range is 0-63.
DSCP mapping for WMM video AC	Select the DSCP used to map WMM video traffic. The supported range is 0-63.
DSCP mapping for WMM best-effort AC	Select the DSCP value used to map WMM best-effort traffic. The supported range is 0-63.
DSCP mapping for WMM background AC	Select the DSCP used to map WMM background traffic. The supported range is 0-63.
Hide SSID	Select this checkbox to enable or disable the hiding of the SSID name in beacon frames. Note that hiding the SSID does very little to increase security.
Deny_Broadcast Probes	When a client sends a broadcast probe request frame to search for all available SSIDs, this option controls whether or not the system responds for this SSID. When enabled, no response is sent and clients have to know the SSID in order to associate to the SSID. When disabled, a probe response frame is sent for this SSID.
Local Probe Request Threshold (dB)	Enter the SNR threshold below which incoming probe requests will get ignored. The supported range of values is 0-100 dB. A value of 0 disables this feature.
Disable Probe Retry	Click this checkbox to enable or disable battery MAC level retries for probe response frames. By default this parameter is enabled, which mean that MAC level retries for probe response frames is disabled.
Battery Boost	Converts multicast traffic to unicast before delivery to the client, thus allowing you to set a longer DTIM interval. The longer interval keeps associated wireless clients from activating their radios for multicast indication and delivery, leaving them in power-save mode longer and thus lengthening battery life. This parameter requires the PEFNG license.
WEP Key 1	First static WEP key associated with the key index. Can be 10 or 26 hex characters in length.

Parameter	Description
WEP Key 2	Second static WEP key associated with the key index. Can be 10 or 26 hex characters in length.
WEP Key 3	Third Static WEP key associated with the key index. Can be 10 or 26 hex characters in length.
WEP Key 4	Fourth Static WEP key associated with the key index. Can be 10 or 26 hex characters in length.
WEP Transmit Key Index	Key index that specifies which static WEP key is to be used. Can be 1, 2, 3, or 4.
WPA Hexkey	WPA pre-shared key (PSK).
WPA Passphrase	WPA passphrase with which to generate a pre-shared key (PSK).
Maximum Transmit Failures	The AP assumes the client has left and should be deauthorized when the AP detects this number of consecutive frames were not delivered because the maximum retry threshold as been exceeded.
BC/MC Rate Optimization	Click this checkbox to enable or disable scanning of all active stations currently associated to an AP to select the lowest transmission rate for broadcast and multicast frames. This option only applies to broadcast and multicast data frames; 802.11 management frames are transmitted at the lowest configured rate. **NOTE:** Do not enable this parameter unless instructed to do so by your Dell technical support representative.
Rate Optimization for delivering EAPOL frames	Click this checkbox to use a more conservative rate for more reliable delivery of EAPOL frames.
Strict Spectralink Voice Protocol (SVP)	Click this checkbox to enable Strict Spectralink Voice Protocol (SVP)
802.11g Beacon Rate	Click this drop-down list to select the beacon rate for 802.11g (use for Distributed Antenna System (DAS) only). Using this parameter in normal operation may cause connectivity problems.
802.11a Beacon Rate	Click this drop-down list to select the beacon rate for 802.11a (use for Distributed Antenna System (DAS) only). Using this parameter in normal operation may cause connectivity problems.
Advertise QBSS Load IE	Click this checkbox to enable the AP to advertise the QBSS load element. The element includes the following parameters that provide information on the traffic situation: • **Station count:** The total number of stations associated to the QBSS. • **Channel utilization:** The percentage of time (normalized to 255) the channel is sensed to be busy. The access point uses either the physical or the virtual carrier sense mechanism to sense a busy channel. • **Available admission capacity:** The remaining amount of medium time (measured as number of 32us/s) available for a station via explicit admission control. The QAP uses these parameters to decide whether to accept an admission control request. A wireless station uses these parameters to choose the appropriate access points. **NOTE:** Ensure that WMM is enabled for legacy APs to advertise the QBSS load element. For 802.11n APs, ensure that either wmm or high throughput is enabled.

Parameter	Description
Advertise Location Information	When this option is enabled, APs broadcast their location within a IE carried in Beacon frames and Probe Response frames. The AP's latitude, longitude and altitude can be configured on the **Configuration > Wireless> AP Installation** page of the controller WebUI, or using the **provision-ap** command in the controller command-line interface.
Advertise AP Name	If this parameter enabled, APs will broadcast the AP name configured by the `ap-name` command. This option is disabled by default.
Enforce User VLAN for Open Stations	Select this option to restrict data traffic from open stations to the user's assigned VLAN. This option is disabled by default.

In the CLI

```
(host)(config) #wlan ssid-profile <profile>
```

WLAN Authentication

The AAA profile configures the authentication for a WLAN. The AAA profile defines the type of authentication (802.1x in this example), the authentication server group, and the default user role for authenticated users.



It is recommended that you assign a unique name to each virtual AP, SSID, and AAA profile that you modify.

Configuring an AAA Profile in the WebUI

1. Navigate to **Configuration > Security > Authentication > Profiles**, the select the **AAA Profiles** tab.
2. Scroll down to the bottom of the **AAA Profiles Summary** pane, then click **Add.** An entry blank appears.
3. Enter the AAA profile name, then click **Add.**
4. In the **profiles list,** and select the AAA profile you just created.
5. Configure the AAA profile parameters (see Table 79),
6. Click **Apply.**

Table 79: *AAA Profile Parameters*

Parameter	Description
Initial role	Click the **Initial Role** drop-down list and select a role for unauthenticated users. The default role for unauthenticated users is **logon**.
MAC Authentication Default Role	Click the MAC Authentication Default Role drop-down list and select the role assigned to the user when the device is MAC authenticated. The default role for MAC authentication is the **guest** user role. If derivation rules are present, the role assigned to the client through these rules take precedence over the default role. **NOTE:** This feature requires the PEFNG license.
802.1X Authentication Default Role	Click the **802.1X Authentication Default Role** drop-down list and select the role assigned to the client after 802.1x authentication. The default role for 802.1x authentication is the **guest** user role. If derivation rules are present, the role assigned to the client through these rules take precedence over the default role. **NOTE:** This feature requires the PEFNG license.
User idle timeout	Select the **Enable** checkbox to configure user idle timeout value for this profile. Specify the idle timeout value for the client in seconds. A value of 0, deletes the user immediately after disassociation from the wireless network. Valid range is 30-15300 in multiples of 30 seconds. Enabling this option overrides the global settings configured in the AAA timers. If this is disabled, the global settings are used.
RADIUS Interim Accounting	When this option is enabled, the RADIUS accounting feature allows the controller to send Interim-Update messages with current user statistics to the server at regular intervals. This option is disabled by default, allowing the controller to send only start and stop messages to the RADIUS accounting server.
User derivation rules	Click the **User derivation** rules drop-down list and specify a user attribute profile from which the user role or VLAN is derived.
Wired to Wireless Roaming	Enable this feature to keep users authenticated when they roam from the wired side of the network. This feature is enabled by default.
SIP authentication role	Click the **SIP authentication** role drop-down list and specify the role assigned to a session initiation protocol (SIP) client upon registration. **NOTE:** This feature requires the PEFNG license.

Parameter	Description
Device Type Classification	When you select this option, the controller will parse user-agent strings and attempt to identify the type of device connecting to the AP. When the device type classification is enabled, the Global client table shown in the **Monitoring>Network > All WLAN Clients** window shows each client's device type, if that client device can be identified.
Enforce DHCP	When you select this option, clients must obtain an IP using DHCP before they are allowed to associate to an AP. Enable this option when you create a user rule that assigns a specific role or VLAN based upon the client device's type. For details, see Working with User-Derived VLANs on page 373. **NOTE:** If a client is removed from the user table by the "Logon user lifetime" AAA timer, then that client will not be able to send traffic until it renews it's DHCP.
PAN firewalls Integration	Requires IP mapping at Palo Alto Networks firewalls. For details, see Palo Alto Networks Firewall Integration on page 620.

7. In the profiles list, select the AAA profile to expand the list of other profiles associated with that AAA prrofile.
8. Click **802.1X Authentication.** The **802.1X Authentication Profile** appears.
   a. Click the **802.1X Authentication Profile** drop-down list and select an authentication profile to associate with your AAA profile.
   b. Click **Apply.**
9. Click **802.1X Authentication Server Group.** The **802.1X Authentication Server Group** appears..
   a. Click the **802.1X Authentication Server Group** drop-down list and select the server group to associate with your AAA profile.
   b. Click **Apply.**
10. Click **MAC Authentication.** The **MAC Authentication Profile** appears.
    a. Click the **MAC Authentication Profile** drop-down list and select a MAC authentication profile to associate with your AAA profile.
    b. Click **Apply.**
11. Click **MAC Authentication Server Group.** The **MAC Authentication Server Group** appears.
    a. Click the **MAC Authentication Server Group** drop-down list and select the MAC server group to associate with your AAA profile.
    b. Click **Apply.**
12. Click **RADIUS Authentication Server Group.** The **RADIUS Authentication Server Group** appears.
    a. Click the **RADIUS Authentication Server Group** drop-down list and select the MAC server group to associate with your AAA profile.
    b. Click **Apply.**

Configuring an AAA Profile in the CLI

```
(host)(config) #aaa authentication dot1x <profile>
(host)(config) #aaa profile <profile>
```

High-Throughput Virtual APs

With the implementation of the IEEE 802.11ac standard, very-high-throughput can be configured to operate on the 5 GHz frequency band. High-throughput (802.11n) can be configured on both the 5 GHz and 2.5 GHz frequency bands. High-throughput is enabled by default, and can be enabled or disabled in the 802.11a and 802.11g radio profiles. For details, see 802.11a and 802.11g RF Management Profiles on page 509

Two different profiles define settings specific to high-throughput APs. The **high-throughput radio profile** defines settings for 40 MHz tolerance, is associated to an AP through it's 802.11a or 802.11g radio profile. The **high-throughput SSID profile** configures the high-throughput SSID settings for 802.11n, and is associated to an AP through it's virtual AP profile

Stations are not allowed to use high-throughput with TKIP standalone encryption, although TKIP can be provided in mixed-mode BSSIDs that support high-throughput. High-throughput is disabled on a BSSID if the encryption mode is standalone TKIP or WEP.



De-aggregation of MAC Service Data Units (A-MSDUs) is supported on the W-3000 Series, W-7220, and the W-6000M3 controllers with a maximum frame transmission size of 4k bytes; however, this feature is always enabled and is not configurable. Aggregation is not currently supported.

Configuring the High-Throughput Radio Profile

You can configure high-throughput radio profile settings using the WebUI or CLI interfaces

In the WebUI

1. Navigate to **Advanced Services > All Profile Management**.
2. In the **Profiles** list, expand the **RF Management** menu, then select **High-throughput radio**.
3. Select an existing profile from the **Profile Details** pane, or enter create a new profile by entering a new name into the entry blank, then clicking **Add**.

   The configuration settings in this profile are divided into two tabs, **Basic** and **Advanced.** The **Basic** tab displays only those configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. If you change a setting on one tab then click and display the other tab without saving your configuration, that setting will revert to its previous value.

Table 80: *High-Throughput Radio Profile Configuration Parameters*

Parameter	Description
Basic	
40MHz intolerance	This parameter controls whether or not APs using this radio profile will advertise intolerance of 40 MHz operation. By default, this option is disabled, and 40 MHz operation is allowed. If you do not want to use 40 Mhz operation, select the **40MHz intolerance** checkbox to enable this feature.
Advanced	

Parameter	Description
honor 40MHz intolerance	When enabled, the radio will stop using the 40 MHz channels if the 40 MHz intolerance indication is received from another AP or station. Uncheck the **Honor 40 Mhz intolerance checkbox** to disable this feature. Default: Enabled
CSD override	Most transmissions to high throughput (HT) stations are sent through multiple antennas using cyclic shift diversity (CSD). When you enable the **CSD Override** parameter, CSD is disabled and only one antenna transmits data, even if they are being sent to high-throughput stations. This enables interoperability for legacy or high-throughput stations that cannot decode 802.11n CDD data. This option is disabled by default, and should only be enabled under the supervision of Dell technical support. Use this feature to turn off antenna diversity when the AP must support legacy clients such as Cisco 7921g VoIP phones, or older 802.11g clients (e.g. Intel Centrino clients). Note, however, that enabling this feature can reduce overall throughput rates.

4. Click **Apply**.

In order for the settings in this profile to take effect, the profile must be associated with an AP's 802.11a or 802.11g radio profile. For details, see the **Associated Profiles** section of , 802.11a/802.11g RF Management Configuration Parameters.

In the CLI

```
(host) (config) rf ht-radio-profile <profile>
  40MHz-intolerance
  clone <profile>
  diversity-spreading-workaround
  honor-40MHz-intolerance
  no...
(host) (config) rf dot11a-radio-profile <profile>
  high-throughput-enable
  ht-radio-profile <profile>
(host) (config) rf dot11g-radio-profile <profile>
  high-throughput-enable
  ht-radio-profile <profile>
```

Configuring the High-Throughput SSID Profile

You can configure high-throughput SSID profile settings using the WebUI or CLI interfaces

In the WebUI

1. Navigate to **Advanced Services > All Profile Management**.
2. In the **Profiles** list, expand the **Wireless LAN** menu, then select **High-throughput radio**.
3. Select an existing profile from the **Profile Details** pane, or enter create a new profile by entering a new name into the entry blank, then clicking **Add**.
4. Configure the high-throughput SSID profile settings described in Table 81 .

   The High-Throughput SSID profile configuration settings are divided into two tabs, **Basic** and **Advanced**. The **Basic** tab displays only those configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. If you change a setting on one tab then click and display the other tab without saving your configuration, that setting will revert to its previous value. Both basic and advanced settings are described in Table 54.

Table 81: *High-Throughput SSID Profile Parameters*

Parameter	Description
Basic High-Throughput SSID Profile Settings	
High throughput enable (SSID)	Determines if this high-throughput SSID allows high-throughput (802.11n) stations to associate. Enabling high-throughput in an WLAN high-throughput SSID profile enables Wi-Fi Multimedia (WMM) base features for the associated SSID.
40 MHz channel usage	Enable or disable the use of 40 MHz channels. This parameter is enabled by default.
Very High throughput enable (SSID)	Enable/Disable support for Very High Throughput (802.11ac) on the SSID.
80 MHz channel usage (VHT)	Enables or disables the use of 80 MHz channels on Very High Throughput (VHT) APs.
VHT - Explicit Transmit Beamforming	Enable or disable VHT Explicit Transmit Beamforming for the W-AP220 Series. When this parameter is enabled, the AP requests information about the Multiple-Input and Multiple-Output (MIMO) channel and uses that information to transmit data over multiple transmit streams using a calculated steering matrix. The result is higher throughput due to improved signal at the beamforming (the receiving client). If this parameter is disabled, all other transmit beamforming settings will not take effect.
Advanced High-Throughput SSID Profile Settings	
VHT - Supported MCS Map	Allows you to set the supported Modulation and Coding Scheme (MCS) map for spatial streams 1 through 3. Each drop down list corresponds to a spatial beginning with 1 on the left and ending with 3 on the right. Default values are set to 9 for each spatial stream.
VHT - Transmit Beamforming Sounding Interval	Time interval in seconds between channel information updates between the AP and the beamformee client. (W-AP220 Series only)
BA AMSDU Enable	Enable/Disable Receive AMSDU in BA negotiation.
Legacy stations	Allow or disallow associations from legacy (non-HT) stations. By default, this parameter is enabled (legacy stations are allowed).
Low-density Parity Check	If enabled, the AP will advertise Low-density Parity Check (LDPC) support. LDPC improves data transmission over radio channels with high levels of background noise.
Maximum number of spatial streams usable for STBC reception	Controls the maximum number of spatial streams usable for STBC reception. 0 disables STBC reception, 1 uses STBC for MCS 0-7. Higher MCS values are not supported. (Supported on the W-AP90 series, W-AP130 Series, W-AP68, W-AP175 and W-AP105 only. The configured value will be adjusted based on AP capabilities.)

Parameter	Description
Maximum number of spatial streams usable for STBC transmission.	Controls the maximum number of spatial streams usable for STBC transmission. 0 disables STBC transmission, 1 uses STBC for MCS 0-7. Higher MCS values are not supported. (Supported on W-AP90 series, W-AP175, W-AP130 Seriesand W-AP105 only. The configured value will be adjusted based on AP capabilities.)
MPDU Aggregation	Enable or disable MAC protocol data unit (MPDU) aggregation. High-throughput APs are able to send aggregated MAC protocol data units (MDPUs), which allow an AP to receive a single block acknowledgment instead of multiple ACK signals. This option, which is enabled by default, reduces network traffic overhead by effectively eliminating the need to initiate a new transfer for every MPDU.
Max received A-MPDU size	Maximum size of a received aggregate MPDU, in bytes. Allowed values: 8191, 16383, 32767, 65535.
Max transmitted A-MPDU size	Maximum size of a transmitted aggregate MPDU, in bytes. Range: 1576–65535
Min MPDU start spacing	Minimum time between the start of adjacent MPDUs within an aggregate MPDU, in microseconds. Allowed values: 0 (No restriction on MDPU start spacing), .25 µsec, .5 µsec, 1 µsec, 2 µsec, 4 µsec.
Short guard interval in 20 MHz mode	Enable or disable use of short (400ns) guard interval in 20 MHz mode. This parameter is enabled by default. A guard interval is a period of time between transmissions that allows reflections from the previous data transmission to settle before an AP transmits data again. An AP identifies any signal content received inside this interval as unwanted inter-symbol interference, and rejects that data. The 802.11n standard specifies two guard intervals: 400ns (short) and 800ns (long). Enabling a short guard interval can decrease network overhead by reducing unnecessary idle time on each AP. Some outdoor deployments, may, however require a longer guard interval. If the short guard interval does not allow enough time for reflections to settle in your mesh deployment, inter-symbol interference values may increase and degrade throughput.

Parameter	Description
Short guard interval in 40 MHz mode	Enable or disable use of short (400ns) guard interval in 40 MHz mode. This parameter is enabled by default. A guard interval is a period of time between transmissions that allows reflections from the previous data transmission to settle before an AP transmits data again. An AP identifies any signal content received inside this interval as unwanted inter-symbol interference, and rejects that data. The 802.11n standard specifies two guard intervals: 400ns (short) and 800ns (long). Enabling a short guard interval can decrease network overhead by reducing unnecessary idle time on each AP. Some outdoor deployments, may, however require a longer guard interval. If the short guard interval does not allow enough time for reflections to settle in your mesh deployment, inter-symbol interference values may increase and degrade throughput.
Supported MCS set	A list of Modulation Coding Scheme (MCS) values or ranges of values to be supported on this SSID. The MCS you choose determines the channel width (20MHz vs. 40MHz) and the number of spatial streams used by the mesh node. The default value is 1–23; the complete set of supported values. To specify a smaller range of values, enter a hyphen between the lower and upper values. To specify a series of different values, separate each value with a comma. Examples: 2–10 1,3,6,9,12 Range: 0–23.
`Temporal Diversity`	When this feature is enabled and the client is not responding to 802.11 packets, the AP will launch two hardware retries; if the hardware retries are not successful then it attempts software retries. This setting is disabled by default.

In order for the settings in this profile to take effect, the profile must be associated with an AP's Virtual AP profile. For details on associating a high-throughput SSID profile with a Virtual AP profile, see Configuring the Virtual AP Profile on page 396

In the CLI

```
(host)(config) wlan ht-ssid-profile <profile-name>
(host)(config) wlan ssid-profile <profile-name>
  ht-ssid-profile <profile>
(host)(config) wlan virtual-ap <profile-name>
  hwlan ssid-profile <profile-name>
```

Guest WLANs

Guest usage in enterprise wireless networks requires the following special consideration:

- Guest users must be separated from employee users by VLANs in the network.
- Guests must be limited not only in where they may go, but also by what network protocols and ports they may use to access resources.
- Guests should be allowed to access only the local resources that are required for IP connectivity. These resources include DHCP and possibly DNS if an outside DNS server is not available. In most cases, a public DNS is always available.

- All other internal resources should be off limits for the guest. This restriction is achieved usually by denying any internal address space to the guest user.
- A time-of-day restriction policy should be used to allow guests to access the network only during normal working hours, because they should be using the network only while conducting official business. A rate limit can also be put on each guest user to keep the user from using up the limited wireless bandwidth.Accounts should be set to expire when their local work is completed, typically at the end of each business day.

The procedures in the following example create an guest WLAN that only allows HTTP and HTTPS traffic from 9:00 a.m. to 5 p.m. on weekdays.

- Configuring a Guest VLAN
- Configuring a Guest Role
- Configuring a Guest Virtual AP

The following sections describe how to do this using the WebUI and the CLI.

Configuring a Guest VLAN

In this example, users on the "Corpnet" WLAN are placed into VLAN 1, which is the default VLAN configured on the controller. For guest users, you need to create another VLAN and assign the VLAN interface an IP address.



Each Virtual AP supports a maximum of 256 VLANs.

In the WebUI

1. Navigate to the **Configuration > Network > VLANs** page.
2. Click **Add** to add a VLAN. Enter 2 in the VLAN ID, and click **Apply**.
3. To assign an IP address and netmask to the VLAN you just created, navigate to the **Configuration >Network > IP > IP Interfaces** page. Click **Edit** for VLAN 2. Enter an IP address and netmask for the VLAN interface, and then click **Apply.**

In the CLI

```
(host)(config) #vlan 2
   interface vlan 2
   ip address <address> <netmask>
```

Configuring a Guest Role

The guest role allows web (HTTP and HTTPS) access only during normal business hours (9:00 a.m. to 5:00 p.m. Monday through Friday).

In the WebUI

1. Navigate to the **Configuration > Security > Access Control > Time Ranges** page.
2. Click **Add**. Enter a name, such as "workhours". Select Periodic. Click **Add**. Under Add Periodic Rule, select Weekday. For Start Time, enter 9:00. For End Time, enter 17:00. Click **Done**. Click **Apply.**
3. Select the **Policies** tab. Click **Add**. Enter a policy name, such as "restricted". From the **Policy Type** drop-down list, select **Session.**
4. Click **Add**.
5. *(Optional)* By default, firewall policies apply to IPv4 clients only. To configure a firewall policy for IPv6 clients, click the **IP Version** drop-down list and select **IPv6.**
6. Click the **Service** drop-down list, select **service,** then select **svc-http**.

7. Click the **Time Range** drop-down list and select the time range you previously configured.
8. Click **Add**.
9. Repeat steps 4-8 to add another rule for the *svc-https* service. Click **Apply.**
10. Select the **User Roles** tab. Click **Add.** Enter guest for Role Name. Under Firewall Policies, click **Add.** Select Choose from Configured Policies and select the policy you previously configured. Click **Done.**
11. Click **Apply**.

In the CLI

```
(host)(config) #time-range workhours periodic
   weekday 09:00 to 17:00
(host)(config) #ip access-list session restricted
   any any svc-http permit time-range workhours
   any any svc-https permit time-range workhours
(host)(config) #user-role guest
   session-acl restricted
```

Configuring a Guest Virtual AP

In this example, you apply the **guest** virtual AP profile to a specific AP.



Best practices are to assign a unique name to each virtual AP, SSID, and AAA profile that you modify. In this example, you use the name guest to identify the virtual AP and SSID profiles.

In the WebUI

1. Navigate to **Configuration > Wireless > AP Configuration > AP Specific** page.
2. Click **New**. Either enter the AP name or select an AP from the list of discovered APs. Click **Add**. The AP name appears in the list.
3. Click **Edit** by the AP name to display the profiles that you can configure for the AP.
4. Expand the **Wireless LAN** profile menu.
5. Select **Virtual AP**.
   a. Click the **Add a profile** drop down list in the **Profile Details** window and select **NEW**.
   b. Enter **guest,** and click **Add**.
   c. Click **Apply**.
6. Click the guest virtual AP to display profile details.
   a. Make sure Virtual AP Enable is selected.
   b. Select 2 for the VLAN.
   c. Click **Apply**.
7. Under Profiles, select the AAA profile under the guest virtual AP profile.
   a. In the Profile Details, select **default-open** from the AAA Profile drop-down list.
   b. Click **Apply.**
8. Under Profiles, select the SSID profile under the guest virtual AP profile.
   a. Select NEW from the SSID Profile drop-down menu.
   b. Enter **guest.**
   c. In the Profile Details, enter **Guest** for the Network Name.
   d. Select None for Network Authentication and Open for Encryption.
   e. Click **Apply**.

In the CLI

```
(host)(config) #wlan ssid-profile guest
   opmode opensystem
(host)(config) #wlan virtual-ap guest
   vap-enable
   vlan 2
   d>eny-time-range workhours
   ssid-profile guest
   aaa-profile default-open
(host)(config) #ap-name building3-lobby
   virtual-ap guest
```

Dell's Adaptive Radio Management (ARM) takes the guesswork out of RF management by using automatic, infrastructure-based controls to maximize client performance and enhance the stability and predictability of the entire Wi-Fi network.

ARM Feature Overviews

The following sections provide a general overview of Adaptive Radio Management feature:

- Understanding ARM on page 432
- Client Match on page 434
- ARM Coverage and Interference Metrics on page 435

Configuring ARM Settings

The section below describes the steps to configure the ARM function to automatically select the best channel and transmission power settings for each AP on your WLAN:

- Configuring ARM Profiles on page 435
- Assigning an ARM Profile to an AP Group on page 443
- Configuring Non-802.11 Noise Interference Immunity on page 450
- Using Multi-Band ARM for 802.11a/802.11g Traffic on page 444
- Reusing Channels to Control RX Sensitivity Tuning on page 449
- Band Steering on page 445
- Enabling Traffic Shaping on page 446
- Spectrum Load Balancing on page 449

ARM Troubleshooting

- Troubleshooting ARM on page 450

Understanding ARM

Dell's Adaptive Radio Management (ARM) technology maximizes WLAN performance even in the highest traffic networks by dynamically and intelligently choosing the best 802.11 channel and transmit power for each Dell AP in its current RF environment.

Dell's ARM technology solves wireless networking challenges such as large deployments, dense deployments, and installations that must support VoIP or mobile users. Deployments with dozens of users per access point can cause network contention and interference, but ARM dynamically monitors and adjusts the network to ensure that all users are allowed ready access. ARM provides the best voice call quality with voice-aware spectrum scanning and call admission control.

With earlier technologies, network administrators would have to perform a site survey at each location to discover areas of RF coverage and interference, and then manually configure each AP according to the results of this survey. Static site surveys can help you choose channel and power assignments for APs, but these surveys are often time-consuming and expensive, and only reflect the state of the network at a single point in time. ARM is more efficient than static calibration, and, unlike older technologies, it continually monitors and adjusts radio resources to provide optimal network performance. Automatic power control can adjust AP

power settings if adjacent APs are added, removed, or moved to a new location within the network, minimizing interference with other WLAN networks. ARM adjusts only the affected APs, so the entire network does not require systemic changes.

ARM Support for 802.11n

ArubaOS version 3.3.x or later supports APs with the 802.11n standard, ensuring seamless integration of 802.11n devices into your RF domain. The Dell AP's 5 Ghz band capacity simplifies the integration of new APs into your legacy network. You can also replace older APs with newer 802.11n-compliant APs while reusing your existing cabling and PoE infrastructure.

A high-throughput (802.11n) AP can use a 40 MHz channel pair comprised of two adjacent 20 MHz channels available in the regulatory domain profile for your country. When ARM is configured for a dual-band AP, it will dynamically select the primary and secondary channels for these devices. It can, however, continue to scan all changes in the a+b/g bands to calculate interference and detect rogue APs.

Monitoring Your Network with ARM

When ARM is enabled, the Dell AP dynamically scans all 802.11 channels in its regulatory domain at regular intervals and will report everything it sees to the controller on each channel it scans. (By default, 802.11n-capable APs scan channels in all regulatory domains.) This includes, but is not limited to, data regarding WLAN coverage, interference, and intrusion detection. You can retrieve this information from the controller to get a quick health check of your WLAN deployment without having to walk around every part of a building with a network analyzer. (For additional information on the individual matrix gathered on the AP's current assigned RF channel, see ARM Coverage and Interference Metrics on page 435.)

Maintaining Channel Quality

Hybrid APs and Spectrum Monitors determine channel quality by measuring channel noise, non-Wi-Fi (interferer) utilization and duty-cycles, and certain types of retries. Regular APs using the ARM feature derive channel quality values by measuring the noise floor for both 802.11 and non-802.11 noise on that channel.

The ARM algorithm is based on what the individual AP hears, so each AP on your WLAN can effectively "self heal" by compensating for changing scenarios like a broken antenna or blocked signals from neighboring APs. Additionally, ARM periodically collects information about neighboring APs to help each AP better adapt to its own changing environment.

Configuring ARM Scanning

The default ARM scanning interval is determined by the **scan-interval** parameter in the ARM profile. If the AP does not have any associated clients (or if most of its clients are inactive), the ARM feature will dynamically readjust this default scan interval, allowing the AP obtain better information about its RF neighborhood by scanning non-home channels more frequently. Starting with ArubaOS 6.2, if an AP attempts to scan a non-home channel but is unsuccessful, the AP will make additional attempts to rescan that channel before skipping it and continuing on to other channels.

The **Over the Air Updates** feature allows an AP to get information about its RF environment from its neighbors, even the AP cannot scan. If you enable this feature, when an AP on the network scans a foreign (non-home) channel, it sends an Over-the-Air (OTA) update in an 802.11 management frame that contains information about that AP's home channel, the current transmission EIRP value of the home channel, and one-hop neighbors seen by that AP.

Starting with ArubaOS 6.3.1, if ARM reports a high noise floor on a channel within a 40 MHz channel pair or 80 MHz channel set, ARM performs an additional 20 MHz scan on each channel within that channel pair or set, to determine the actual noise floor of each affected channel. This allows ARM to avoid assigning the overused channel, while still allowing channel assignments to the other unaffected channels in that channel pair or set.

Understanding ARM Application Awareness

Dell APs keep a count of the number of data bytes transmitted and received by their radios to calculate the traffic load. When a WLAN gets very busy and traffic exceeds a predefined threshold, load-aware ARM dynamically adjusts scanning behavior to maintain uninterrupted data transfer on heavily loaded systems. ARM-enabled APs will resume their complete monitoring scans when the traffic has dropped to normal levels. You can also define a firewall policy that pauses ARM scanning when the AP detects critically important or latency-sensitive traffic from a specified host or network.

ARM's band steering feature encourages dual-band capable clients to stay on the 5GHz band on dual-band APs. This frees up resources on the 2.4GHz band for single band clients like VoIP phones.

The ARM "Mode Aware" option is a useful feature for single radio, dual-band WLAN networks with high density AP deployments. If there is too much AP coverage, those APs can cause interference and negatively impact your network. Mode aware ARM can turn APs into Air Monitors if necessary, then turn those Air Monitors back into APs when they detect gaps in coverage. Note that an Air Monitor will not turn back into an AP if it detects client traffic (or client traffic increases), but will change to an AP only if it detects coverage holes.

Client Match

The ARM client match feature continually monitors a client's RF neighborhood to provide ongoing client bandsteering and load balancing, and enhanced AP reassignment for roaming mobile clients. This feature is recommended over the legacy bandsteering and spectrum load balancing features, which, unlike client match, do not trigger AP changes for clients already associated to an AP.



Legacy 802.11a/b/g devices do not support the client match feature. When you enable client match on 802.11n-capable devices, the client match feature overrides any settings configured for the legacy bandsteering, station handoff assist or load balancing features. 802.11ac-capable devices do not support the legacy bandsteering, station hand off or load balancing settings, so these APs must be managed on using client match.

When you enable this feature on an AP, that AP is responsible for measuring the RF health of its associated clients. The AP receives and collects information about clients in its neighborhood, and periodically sends this information to the controller. The controller aggregates information it receives from all APs using client match, and maintains information for all associated clients in a database. The controller shares this database with the APs (for their associated clients), and the APs use the information to compute the client-based RF neighborhood and determine which APs should be considered candidate APs for each client. When the controller receives a client steer request from an AP, the controller identifies the optimal AP candidate and manages the client's relocation to the desired radio. This is an improvement from previous releases, where the ARM feature was managed exclusively by APs, the without the larger perspective of the client's RF neighborhood.

The following client/AP mismatch conditions are managed by the client match feature:

- **Load Balancing**: Client match balances clients across APs on different channels, based upon the client load on the APs and the SNR levels the client detects from an underused AP. If an AP radio can support additional clients, the AP will participate in client match load balancing and clients can be directed to that AP radio, subject to predefined SNR thresholds.
- **Sticky Clients**: The client match feature also helps mobile clients that tend to stay associated to an AP despite low signal levels. APs using client match continually monitor the client's RSSI as it roams between APs, and move the client to an AP when a better radio match can be found. This prevents mobile clients from remaining associated to an APs with less than ideal RSSI, which can cause poor connectivity and reduce performance for other clients associated with that AP.
- **Band Steering/Band Balancing**: APs using the client match feature monitor the RSSI for clients that advertise a dual-band capability. If a client is currently associated to a 2.4 GHz radio and the AP detects that

the client has a good RSSI from the 5 Ghz radio, the controller will attempt to steer the client to the 5 Ghz radio, as long as the 5 Ghz RSSI is not significantly worse than the 2.4 GHz RSSI, and the AP retains a suitable distribution of clients on each of its radios.



The client match feature is enabled through the AP's ARM profile. Although default client match settings are recommended for most users, advanced client match settings can be configured using **rf arm-profile** commands in the command-line interface.

ARM Coverage and Interference Metrics

ARM computes coverage and interference metrics for each valid channel, and chooses the best performing channel and transmit power settings for each AP's RF environment. Each AP gathers other metrics on their ARM-assigned channel to provide a snapshot of the current RF health state.

The following two metrics help the AP decide which channel and transmit power setting is best.

- **Coverage Index:** The AP uses this metric to measure RF coverage. The coverage index is calculated as x/y, where "x" is the AP's weighted calculation of the Signal-to-Noise Ratio (SNR) on all valid APs on a specified 802.11 channel, and "y" is the weighted calculation of the Dell AP's SNR the neighboring APs see on that channel.

  To view these values for an AP in your current WLAN environment, issue the CLI command **show ap arm rf-summary ap-name <ap-name>**, where **<ap-name>** is the name of an AP for which you want to view information.
- **Interference Index:** The AP uses this metric to measure co-channel and adjacent channel interference. The Interference Index is calculated as a/b//c/d, where:
  - Metric value "a" is the channel interference the AP sees on its selected channel.
  - Metric value "b" is the interference the AP sees on the adjacent channel.
  - Metric value "c" is the channel interference the AP's neighbors see on the selected channel.
  - Metric value "d" is the interference the AP's neighbors see on the adjacent channel.

  To manually calculate the total Interference Index for a channel, issue the CLI command **show ap arm rf-summary ap-name <ap-name>**, then add the values *a+b+c+d.*

Each AP also gathers the following additional metrics, which can provide a snapshot of the current RF health state. View these values for each AP using the CLI command **show ap arm rf-summary ip-addr <ap ip address>**.

- Amount of Retry frames (measured in %)
- Amount of Low-speed frames (measured in %)
- Amount of Non-unicast frames (measured in %)
- Amount of Fragmented frames (measured in %)
- Amount of Bandwidth seen on the channel (measured in kbps)
- Amount of PHY errors seen on the channel (measured in %)
- Amount of MAC errors seen on the channel (measured in %)
- Noise floor value for the specified AP

Configuring ARM Profiles

ARM profile settings are divided into two categories: **Basic** and **Advanced.** The Basic ARM settings include ARM scanning checkbox and general configuration parameters such as channel and power assignments and minimum and maximum allowed EIRP values. Most network environments do not require any changes to the

advanced ARM configuration settings. If, however, your network supports a large amount of VoIP or Video traffic, or if you have unusually high security requirements you may want to manually adjust the basic ARM thresholds.

Creating and Configuring a New ARM Profile

There are two ways to create a new ARM profile. You can create an entirely new profile with all default settings using the WebUI or CLI interfaces, or you can make a copy of an existing profile using the CLI interface.

In the WebUI

To create a new ARM profile with all default settings via the WebUI:

1. Select **Configuration > Advanced Services > All Profiles**. The **All Profile Management** window opens.
2. Select **RF Management** to expand the **RF Management** section.
3. Select **Adaptive Radio Management (ARM) Profile**. Any currently defined ARM profiles appears in the right pane of the window. If you have not yet created any ARM profiles, this pane displays the **default** profile only.
4. To create a new profile with all default settings, enter a name in the entry blank. The name must be 1–63 characters, and can be composed of alphanumeric characters, special characters, and spaces. If your profile name includes a space, it must be enclosed within quotation marks.
5. Click **Add**. The new profile appears in the ARM profile list.
6. Select the name of that profile to display the current configuration settings of that profile.

To create a new ARM profile via the command-line interface, access the CLI in config mode and issue the following command:

```
(host)(config) #rf arm-profile <profile>
```

The name must be 1–63 characters, and can be composed of alphanumeric characters, special characters and spaces. If your profile name includes a space, it must be enclosed within quotation marks

Table 82: *ARM Profile Configuration Parameters*

Setting	Description
Basic Configuration Settings	
Assignment	Activates one of four ARM channel/power assignment modes. (The default value is **single-band**.) • disable: Disables ARM calibration and reverts APs back to default channel and power settings specified by the AP's radio profile. • maintain: APs maintain their current channel and power settings. This setting can be used to maintain AP channel and power levels after ARM has initially selected the best settings. • multi-band: For single-radio APs, this value computes ARM assignments for both 5 GHZ (802.11a) and 2.4 GHZ (802.11b/g) frequency bands. • single-band: For dual-radio APs, this value enables APs to change transmit power and channels within their same frequency band, and to adapt to changing channel conditions.
Allowed bands for 40MHz channels	The specified setting allows ARM to determine if 40 MHz mode of operation is allowed on the 5 GHz or 2.4 GHz frequency band only, on both frequency bands, or on neither frequency band.
80MHz support	If enabled, this feature allows ARM to assign 80 MHz channels on APs that support VHT. This setting is enabled by default.

Table 82: *ARM Profile Configuration Parameters*

Setting	Description
Max Tx EIRP	Maximum effective isotropic radiated power (EIRP) from 3 to 33 dBm in 3 dBm increments. You may also specify a special value of 127 dBm for regulatory maximum. Higher power level settings may be constrained by local regulatory requirements and AP capabilities. In the event that an AP is configured for a **Max Tx EIRP** setting it cannot support, this value will be reduced to the highest supported power setting. The default value for this parameter is 127 dBm. **NOTE:** Power settings will not change if the **Assignment** option is set to **disabled** or **maintain**.
Min Tx EIRP	Minimum effective isotropic radiated power (EIRP) from 3 to 33 dBm in 3 dBm increments. You may also specify a special value of 127 dBm for regulatory maximum to disable power adjustments for environments such as outdoor mesh links. Note that power settings will not change if the **Assignment** option is set to **disabled** or **maintain**. Higher power level settings may be constrained by local regulatory requirements and AP capabilities. In the event that an AP is configured for a **Min Tx EIRP** setting it cannot support, this value will be reduced to the highest supported power setting. The default value for this parameter is 9 dBm. Consider configuring a **Min Tx Power** setting higher than the default value if most of your APs are placed on the ceiling. APs on a ceiling often have good line of sight between them, which will cause ARM to decrease their power to prevent interference. However, if the wireless clients down on the floor do not have such a clear line back to the AP, you could end up with coverage gaps.
Client Match	The client match feature helps optimize network resources by balancing clients across channels, regardless of whether the AP or the controller is responding to the wireless clients' probe requests. If enabled, the controller compares whether an AP has more clients than its neighboring APs on other channels. If an AP's client load is at or over a predetermined threshold as compared to its immediate neighbors, or if a neighboring Dell AP on another channel does not have any clients, load balancing will be enabled on that AP. This feature is enabled by default. For details, see Client Match on page 434.
Scanning	The **Scanning** checkbox enables or disables AP scanning across multiple channels. This checkbox is selected by default. Do not disable scanning unless you want to disable ARM and manually configure AP channel and transmission power. Disabling this option also disables the following scanning features: • Multi Band Scan • Rogue AP Aware • Voip Aware Scan • Power Save Scan
Multi Band Scan	If enabled, single radio channel APs scans for rogue APs across multiple channels. This option requires that **Scanning** is also enabled. (The **Multi Band Scan** option does not apply to APs that have two radios, as these devices already scan across multiple channels. If one of these dual-radio devices are assigned an ARM profile with Multi Band enabled, that device will ignore this setting.) Default: disabled
VoIP Aware Scan	Dell's VoIP Call Admission Control (CAC) prevents any single AP from becoming congested with voice calls. When you enable CAC, you should also enable **VoIP Aware Scan** in the ARM profile, so the AP will not attempt to scan a different channel if one of its clients has an active VoIP call. This option requires that **Scanning** is also enabled. Default: disabled
Power Save Aware Scan	If enabled, the AP will not scan a different channel if it has one or more clients that is in power save mode. Default: disabled

Table 82: *ARM Profile Configuration Parameters*

Setting	Description
Video Aware Scan	As long as there is at least one video frame every 100 mSec the AP will reject an ARM scanning request. Note that for each radio interface, video frames must be defined in one of two ways: • Classify the frame as video traffic via a session ACL. • Enable WMM on the WLAN's SSID profile and define a specific DSCP value as a video stream. Next, create a session ACL to tag the video traffic with the that DSCP value.
Scan Mode	By default, 802.11n-capable APs scan channels within all regulatory domains. To limit the AP scans to just the regulatory domain for that AP, click the **Scan Mode** drop-down list and select **reg-domain**. **NOTE:** This setting does not apply to APs that do not support 802.11n; these APs will scan their regulatory domain only.
Client Match	Select this checkbox to enable the client match feature, which monitors clients' RF neighborhood to provide ongoing client bandsteering and load balancing, and enhanced AP reassignment for roaming mobile clients. For complete information on this feature, see Client Match on page 434.
Advanced Configuration Settings	
Assignment	Activates one of four ARM channel/power assignment modes: • disable: Disables ARM calibration and reverts APs back to default channel and power settings specified by the AP's radio profile. • maintain: APs maintain their current channel and power settings. This setting can be used to maintain AP channel and power levels after ARM has initially selected the best settings. • multi-band: For single-radio APs, this value computes ARM assignments for both 5 GHZ (802.11a) and 2.4 GHZ (802.11b/g) frequency bands. • single-band: For dual-radio APs, this value enables APs to change transmit power and channels within their same frequency band, and to adapt to changing channel conditions. Default: single-band
Allowed bands for 40MHz channels	The specified setting allows ARM to determine if 40 MHz mode of operation is allowed on the 5 GHz or 2.4 GHz frequency band only, on both frequency bands, or on neither frequency band.
Client Aware	If the **Client Aware** option is enabled, the AP does not change channels if there is an active client associated to that AP. (Activity is defined by the **sta-inactivity-time** parameter in the IDS general profile. By default, a client is considered active if it has sent or received traffic within the last 60 seconds.) If you disable **Client Aware** , the AP may change to a more optimal channel, but this change may also disrupt current client traffic. Default: enabled
Max Tx EIRP	Maximum effective isotropic radiated power (EIRP) from 3 to 33 dBm in 3 dBm increments. You may also specify a special value of 127 dBm for regulatory maximum. Higher power level settings may be constrained by local regulatory requirements and AP capabilities. In the event that an AP is configured for a **Max Tx EIRP** setting it cannot support, this value will be reduced to the highest supported power setting. Default: 127 dBm **NOTE:** Power settings will not change if the **Assignment** option is set to **disabled** or **maintain**.

Table 82: *ARM Profile Configuration Parameters*

Setting	Description
Min Tx EIRP	Minimum effective isotropic radiated power (EIRP) from 3 to 33 dBm in 3 dBm increments. You may also specify a special value of 127 dBm for regulatory maximum to disable power adjustments for environments such as outdoor mesh links. Note that power settings will not change if the **Assignment** option is set to **disabled** or **maintain**. Higher power level settings may be constrained by local regulatory requirements and AP capabilities. In the event that an AP is configured for a **Min Tx EIRP** setting it cannot support, this value will be reduced to the highest supported power setting. Default: 9 dBm **NOTE:** Consider configuring a **Min Tx Power** setting higher than the default value if most of your APs are on the ceiling. APs on a ceiling often have good line of sight between them, which will cause ARM to decrease their power to prevent interference. However, if the wireless clients down on the floor do not have such a clear line back to the AP, you could end up with coverage gaps.
Rogue AP Aware	If you have enabled both the **Scanning** and **Rogue AP options**, Dell APs may change channels to contain off-channel rogue APs with active clients. This security features allows APs to change channels even if the **Client Aware** setting is disabled. This setting is disabled by default, and should only be enabled in high-security environments where security requirements are allowed to consume higher levels of network resources. You may prefer to receive Rogue AP alerts via SNMP traps or syslog events. Default: disabled
Scan Interval	If **Scanning** is enabled, the **Scan Interval** defines how often the AP will leave its current channel to scan other channels in the band. Off-channel scanning can impact client performance. Typically, the shorter the scan interval, the higher the impact on performance. If you are deploying a large number of new APs on the network, you may want to lower the Scan Interval to help those APs find their optimal settings more quickly. Raise the Scan Interval back to its default setting after the APs are functioning as desired. Range: 0–2,147,483,647 seconds. Default: 10 seconds
Active Scan	When you enable **Active Scan**, an AP initiates active scanning via probe request. This option elicits more information from nearby APs, but also creates additional management traffic on the network. **Active Scan** is disabled by default, and should *not be enabled* except under the direct supervision of Dell Support. Default: disabled
ARM Over the Air Updates	The **ARM Over the Air Updates** option allows an AP to get information about its RF environment from its neighbors, even the AP cannot scan. If this feature is enabled, when an AP on the network scans a foreign (non-home) channel, it sends other APs an Over-the-Air (OTA) update in an 802.11 management frame that contains information about the scanning AP's home channel, the current transmission EIRP value of its home channel, and one-hop neighbors seen by that AP. Default: enabled
Scanning	The **Scanning** checkbox enables or disables AP scanning across multiple channels. Disabling this option also disables the following scanning features: • Multi Band Scan • Rogue AP Aware • Voip Aware Scan • Power Save Scan Do not disable Scanning unless you want to disable ARM and manually configure AP channel and transmission power. Default: enabled

Table 82: *ARM Profile Configuration Parameters*

Setting	Description
Multi Band Scan	If enabled, single radio channel APs scans for rogue APs across multiple channels. This option requires that **Scanning** is also enabled. (The **Multi Band Scan** option does not apply to APs that have two radios, as these devices already scan across multiple channels. If one of these dual-radio devices are assigned an ARM profile with Multi Band enabled, that device will ignore this setting.) Default: disabled
VoIP Aware Scan	Dell's VoIP Call Admission Control (CAC) prevents any single AP from becoming congested with voice calls. When you enable CAC, you should also enable **VoIP Aware Scan** in the ARM profile, so the AP will not attempt to scan a different channel if one of its clients has an active VoIP call. This option requires that **Scanning** is also enabled. Default: disabled
Power Save Aware Scan	If enabled, the AP will not scan a different channel if it has one or more clients that is in power save mode. Default: disabled
Video Aware Scan	As long as there is at least one video frame every 100 mSec the AP will reject an ARM scanning request. Note that for each radio interface, video frames must be defined in one of two ways: • Classify the frame as video traffic via a session ACL. • Enable WMM on the WLAN's SSID profile and define a specific DSCP value as a video stream. Next, create a session ACL to tag the video traffic with the that DSCP value.
Ideal Coverage Index	The Dell coverage index metric is a weighted calculation based on the RF coverage for all Dell APs and neighboring APs on a specified channel. The **Ideal Coverage Index** specifies the ideal coverage that an AP should try to achieve on its channel. The denser the AP deployment, the lower this value should be. Range: 2–20 Default: 10 For additional information on how this the Coverage Index is calculated, see ARM Coverage and Interference Metrics on page 435
Acceptable Coverage Index	For multi-band implementations, the **Acceptable Coverage Index** specifies the minimal coverage an AP it should achieve on its channel. The denser the AP deployment, the lower this value should be. Range: 1–6 Default: 4
Free Channel Index	The Dell Interference index metric measures interference for a specified channel and its surrounding channels. This value is calculated and weighted for all APs on those channels (including 3rd-party APs). An AP will only move to a new channel if the new channel has a lower interference index value than the current channel. **Free Channel Index** specifies the required difference between the two interference index values before the AP moves to the new channel. The lower this value, the more likely it is that the AP will move to the new channel. The range of possible values is 10–40. Default: 25 For additional information on how this the Channel Index is calculated, see ARM Coverage and Interference Metrics on page 435.
Backoff Time	After an AP changes channel or power settings, it waits for the backoff time interval before it asks for a new channel/power setting. Range: 120–3600 seconds. Default: 240 seconds

Table 82: *ARM Profile Configuration Parameters*

Setting	Description
Error Rate Threshold	The minimum percentage of PHY errors and MAC errors in the channel that will trigger a channel change. Default: 50%
Error Rate Wait Time	Minimum time in seconds the error rate has to exceed the Error Rate Threshold before it triggers a channel change. Default: 30 seconds
Channel Quality Aware Arm	Select this checkbox to allow ARM to initiate a channel change due to low quality on the current channel.
Channel Quality Threshold	Channel quality percentage below which ARM initiates a channel change. Range: 0-100% Default: 70%.
Channel Quality Wait Time	If channel quality is below the specified channel quality threshold for this wait time period, ARM initiates a channel change. Range:1-3600 seconds Default: 120 seconds.
Minimum Scan Time	Minimum number of times a channel must be scanned before it is considered for assignment. Range: 0–2,147,483,647 scans. It is recommended to use a **Minimum Scan Time** between 1–20 scans. Default: 8 scans
Load Aware Scan Threshold	Load aware ARM preserves network resources during periods of high traffic by temporarily halting ARM scanning if the load for the AP gets too high. The **Load Aware Scan Threshold** is the traffic throughput level an AP must reach before it stops scanning. Range: 0–20,000,000 bytes/second. (Specify 0 to disable this feature.) Default: 1250000 Bps
Mode Aware ARM	If enabled, ARM will turn APs into Air Monitors (AMs) if it detects higher coverage levels than necessary. This helps avoid higher levels of interference on the WLAN. Although this setting is disabled by default, you may want to enable this feature if your APs are deployed in close proximity (less than 60 feet apart). Mode aware ARM turns Air Monitors back into APs when they detect gaps in coverage. Note that an Air Monitor will not turn back into an AP if it detects client traffic (or client traffic increases), but will change to an AP only if it detects coverage holes. Default: disabled
Scan Mode	By default, 802.11n-capable APs scan channels within all regulatory domains. To limit the AP scans to just the regulatory domain for that AP, click the **Scan Mode** drop-down list and select **reg-domain.** **NOTE:** This setting does not apply to APs that do not support 802.11n; these APs will scan their regulatory domain only.
Video Aware Scan	As long as there is at least one video frame every 100 mSec the AP will reject an ARM scanning request. Note that for each radio interface, video frames must be defined in one of two ways: • Classify the frame as video traffic via a session ACL. • Enable WMM on the WLAN's SSID profile and define a specific DSCP value as a video stream. Next, create a session ACL to tag the video traffic with the that DSCP value.

In the CLI

You must be in config mode to create, modify or delete an ARM profile using the CLI. Specify an existing ARM profile with the <profile-name> parameter to modify an existing ARM profile, or enter a new name to create an entirely new profile.

Configuration details and any default values for each of these parameters are described in Table 82. If you do not specify a parameter for a new profile, that profile uses the default value for that parameter. Put the **no** option before any parameter to remove the current value for that parameter and return it to its default setting.



The ARM profile includes advanced client match settings that can be configured through the command-line interface only. The default values for these settings are recommended for most users, and caution should be used when changing them to a non-default value. For complete details on all client match configuration settings, refer to the *Dell Networking W-Series ArubaOS CLI Reference Guide*.

Use the following commands to create or modify an ARM profile:

```
(host)(config) #rf arm-profile <profile>
rf arm-profile <profile>
   40MHz-allowed-bands {All|None|a-only|g-only}
   80MHz support
   acceptable-coverage-index <number>
   active-scan (not intended for use)
   assignment {disable|maintain|multi-band|single-band}
   backoff-time <seconds>
   channel-quality-aware-arm
   channel-quality-threshold <channel-quality-threshold>
   channel-quality-wait-time <seconds>
   client-aware
   client-match
   clone <profile>
   cm-blist-timeout <secs>
   cm-lb-client-thresh <#-of-clients>
   cm-lb-snr-thresh <dB>
   cm-lb-thresh <%-of-clients>
   cm-max-steer-fails <#-of-fails>
   cm-stale-age <secs>
   cm-sticky-check_intvl <secs>
   cm-sticky-check_snr <dB>
   cm-sticky-min-signal <-dB>
   cm-sticky-snr-thresh
   cm-update-interval <dB>
   error-rate-threshold <percent>
   error-rate-wait-time <seconds>
   free-channel-index <number>
   ideal-coverage-index <number>
   load-aware-scan-threshold
   max-tx-power <dBm>
   min-scan-time <# of scans>
   min-tx-power <dBm>
   mode-aware
   multi-band-scan
   no ...
   ota-updates
   ps-aware-scan
   rogue-ap-aware
   scan mode all-reg-domain|reg-domain
   scanning
   video-aware-scan
   voip-aware-scan
```

Modifying an Existing Profile

To modify an existing ARM profile:

1. Follow steps 1–3 in the above procedure to access the **Adaptive Radio Management (ARM) profile** window.
2. From the list of profiles, select the profile with the settings you would like to modify.
3. Make any desired changes to the parameters described in Table 82, then click **Apply**.

To modify of an existing ARM profile via the command-line interface, access the CLI in config mode and issue the following command.

```
(host)(config) #rf arm-profile <profile>
```

Copying an Existing Profile

To create a new ARM profile based upon the settings of another existing profile:

1. Follow steps 1–3 in the above procedure to access the **Adaptive Radio Management (ARM) profile** window.
2. From the list of profiles, select the profile with the settings you would like to copy.
3. Click **Save As**.
4. Enter a name for the new profile in the entry blank. The name must be 1–63 characters, and can be composed of alphanumeric characters, special characters and spaces.
5. Click **Apply**.

To create a copy of an existing ARM profile via the command-line interface, access the CLI in config mode and issue the following command, where **<newprofile>** is a unique name for the new ARM profile, and **<profile>** is the name of the existing profile whose setting you want to copy.

```
(host)(config) #rf arm-profile <newprofile> clone <profile>
```

Deleting a Profile

You can only delete unused ARM profiles; Dell will not let you delete an ARM profile that is currently assigned to an AP group.

To delete an ARM profile In the WebUI:

1. Select **Configuration > Advanced Services > All Profiles**. The **All Profile Management** window opens.
2. Select **RF Management** to expand the **RF Management** section.
3. Select **Adaptive Radio Management (ARM) Profile**.
4. Select the name of the profile you want to delete.
5. Click **Delete**.

To delete an ARM profile using the CLI, issue the following command where **<profile>** is the name of the ARM profile you want to remove:

```
(host)(config) #no rf arm-profile <profile>
```

Assigning an ARM Profile to an AP Group

Once you have created a new ARM profile, you must assign it to a group of APs before those ARM settings go into effect. Each AP group has a separate set of configuration settings for its 802.11a radio profile and its 802.11g radio profile. You can assign the same ARM profile to each radio profile, or select different ARM profiles for each radio.

In the WebUI

To assign an ARM profile to an AP group via the Web User Interface:

1. Select **Configuration > Wireless > AP Configuration.**
2. Click the **AP Group** tab if it is not already selected.
3. Click the **Edit** button beside the AP group to which you want to assign the new ARM profile.
4. Expand the **RF Management** section in the left window pane.
5. Select a radio profile for the new ARM profile.
   - To assign a new ARM profile to an AP group's 802.11a radio profile, expand **the 802.11a radio profile** section.
   - To assign a new ARM profile to an AP group's 802.11g radio profile, expand the **802.11g radio profile** section.
6. Select **Adaptive Radio management (ARM) Profile**.
7. Click the **Adaptive Radio Management (ARM) Profile** drop-down list in the right window pane, and select a new ARM profile.
8. (Optional) repeat steps 6–8 to assign an ARM profile to another 802.11a or 802.11g radio profile.
9. Click **Apply**.

You can also assign an ARM profile to an AP group by selecting a radio profile, identifying an AP group assigned to that radio profile, and then assigning an ARM profile to one of those groups.

1. Select **Configuration > Advanced Services> All Profiles**.
2. Select **RF Management**, and then expand either the **802.11a radio profile** or **802.11b radio profile**.
3. Select an individual radio profile name to expand that profile.
4. Click **Adaptive Radio Management (ARM) Profile**, and then use the **Adaptive Radio management (ARM) Profile** drop-down list in the right window pane to select a new ARM profile for that radio.

In the CLI

To assign an ARM profile to an AP group via the command-line interface, access the CLI in config mode and issue the following commands where **<ap_profile>** is the name of the AP group, and **<arm_profile>** is the name of the ARM profile you want to assign to that radio band.:

```
(host)(config) #rf dot11a-radio-profile <ap_profile>
  arm-profile <arm_profile>
```

and

```
(host)(config) #rf dot11g-radio-profile <ap_profile>
  arm-profile <arm_profile>
```

Using Multi-Band ARM for 802.11a/802.11g Traffic

It is recommended that you use the **multi-band** ARM assignment and **Mode Aware** ARM feature for single-radio APs in networks with traffic in the 802.11a and 802.11g bands. This feature allows a single-radio AP to dynamically change its radio bands based on current coverage on the configured band. This feature is enabled via the AP's ARM profile.

When you first provision a single-radio AP, it initially operates in the radio band specified in its AP system profile. If the AP finds adequate coverage on multiple channels in its current band of operation, the **mode-aware** feature allows the AP to temporarily turn itself off and become an AP Air Monitor (APM). In AP Monitor mode, the AP scans all channels across both bands to verify that each channel meets or exceeds its required level of acceptable radio coverage (as defined by the in the ARM profile).

If the AP Monitor detects that a channel on the 802.11g band does not have adequate radio coverage, it will convert back to an AP on that 802.11 channel. If the 802.11g band is adequately covered, the AP Monitor will next check the 802.11a band. If a channel on the 802.11a band lacks coverage, the AP Monitor will convert back to an AP on that 802.11a channel.

Band Steering

ARM's band steering feature encourages dual-band capable clients to stay on the 5GHz band on dual-band APs, freeing up resources on the 2.4GHz band for single-band clients like VoIP phones.Band steering reduces co-channel interference and increases available bandwidth for dual-band clients, because there are more channels on the 5GHz band than on the 2.4GHz band. Dual-band 802.11n-capable clients may see even greater bandwidth improvements, because the band steering feature will automatically select between 40MHz or 20MHz channels in 802.11n networks. This feature is disabled by default, and must be enabled in a Virtual AP profile.

The band steering feature considers several metrics before it determines if a client should be steered to the 5GHz band, including client RSSI. For example, this feature will only steer a client to the 5GHz band if that client detects an acceptable RSSI value from an 5GHz AP radio, and the signal from the 5Ghz radio is not significantly weaker than the RSSI from the 2.4GHz radio.

This feature also takes into account the current load on each radio of a dual-band AP. The band steering feature will *not* steer more clients to 5G on that AP if there are many clients associated to the AP, and significantly more 802.11a clients than 80211g clients.b

The band steering feature supports both campus APs and remote APs that have a virtual AP profile set to tunnel, split-tunnel, or bridge forwarding mode. Note, however, that if a campus or remote AP has virtual AP profiles configured in bridge or split-tunnel forwarding mode *but no virtual AP in tunnel mode*, those APs will gather information about 5G-capable clients independently and will not exchange this information with other APs that also have bridge or split-tunnel virtual APs only. The band steering feature will not proactively disconnect clients that are already associated with a radio. All band steering occurs when a client is trying to associate to a new AP radio.



Best practices are to use either the Band Steering or the Client Match feature to balance client loads, but not both at the same time.

Steering Modes

Band steering supports the following three different band steering modes:

- **Prefer-5GHz** *(Default)*: If you configure the AP to use **prefer-5GHz** band steering mode, the AP will not respond to 2.4 Ghz probe requests from a client if all the following conditions are met.
  - The client has already probed the AP on the 5Ghz band and therefore is known to be capable of sending probes on the 5Ghz band.
  - The client is not currently associated on the 2.4Ghz radio to this AP.
  - The client has sent fewer than 8 probes in the last 10 seconds. If the client has sent more than 8 probes in the last 10 seconds, the client will be able to connect using whatever band it prefers
- **Force-5GHz:** When the AP is configured in **force-5GHz** band steering mode, the AP will not respond to 2.4 Ghz probe requests from a client if all the following conditions are met.
  - The client has already probed the AP on the 5Ghz band and therefore is known to be capable of sending probes on the 5Ghz band.
  - The client is not currently associated on the 2.4Ghz radio of this AP.
- **Balance-bands:** In this band steering mode, the AP uses client load and RSSI information balance the clients across the two radios and best use the available 2.4G bandwidth. This feature takes into account the

fact that the 5Ghz band has more channels than the 2.4 Ghz band, and that the 5Ghz channels operate in 40MHz while the 2.4Ghz band operates in 20MHz.

Enabling Band Steering

Band steering is configured in a virtual AP profile. Use the following procedures to enable or disable Band Steering using the WebUI or command-line interfaces.

In the WebUI

1. Select **Configuration > Advanced Services > All Profiles.** The **All Profile Management** window opens.
2. Select **Wireless LAN** to expand the **Wireless LAN** section.
3. Select **Virtual AP profile** to expand the **Virtual AP Profile** section.
4. Select the name of the Virtual AP profile for which you want to enable band steering.

   (To create a new virtual AP profile, enter a name for a new profile in the **Profile Details** window, then click **Add**. The new profile will appear in the **Profiles** list. Select that profile to open the **Profile Details** pane.)
5. In the **Profile Details** pane, select **Band Steering** to enable this feature, or uncheck the **Band Steering** checkbox to disable this feature.
6. Once band steering is enabled, click the **steering mode** drop-down list and select the desired steering mode.
7. Click **Apply.**

In the CLI

Use the following commands to enable band steering via the command-line interface. Access the CLI in config mode then specify an existing virtual AP with the <name> parameter to modify an existing profile, or enter a new name to create an entirely new virtual AP profile:

```
(host)(config) #wlan virtual-ap <profile> band-steering
(host)(config) #wlan virtual-ap <profile> steering-mode balance-bands|force-5ghz|prefer-5ghz
```

To disable band steering, include the **no** parameter :

```
(host)(config) #wlan virtual-ap <profile> no band-steering
```

You can also use the command-line interface to configure and apply multiple instances of virtual AP profiles to an AP group or to an individual AP. Use the following commands to apply a virtual AP profile to an AP group or an individual AP:

```
(host)(config) #ap-group <name> virtual-ap <profile>
(host)(config) #ap-name <name> virtual-ap <profile>
```

Enabling Traffic Shaping

In a mixed-client network, it is possible for slower clients to bring down the performance of the whole network. To solve this problem and ensure fair access to all clients independent of their WLAN or IP stack capabilities, an AP can implement the traffic shaping feature. This feature has the following three options:

- **default-access:** Traffic shaping is disabled, and client performance is dependent on MAC contention resolution. This is the default traffic shaping setting.
- **fair-access:** Each client gets the same airtime, regardless of client capability and capacity. This option is useful in environments like a training facility or exam hall, where a mix of 802.11a/g, 802.11g and 802.11n clients need equal to network resources, regardless of their capabilities.
- **preferred-access:** High-throughput (802.11n) clients do not get penalized because of slower 802.11a/g or 802.11b transmissions that take more air time due to lower rates. Similarly, faster 802.11a/g clients get more access than 802.11b clients.

With this feature, an AP keeps track of all BSSIDs active on a radio, all clients connected to the BSSID, and 802.11a/g, 802.11b, or 802.11n capabilities of each client. Every sampling period, airtime is allocated to each client, giving it opportunity to get and receive traffic. The specific amount of airtime given to an individual client is determined by the following factors:

- Client capabilities (802.11a/g, 802.11b or 802.11n).
- Amount of time the client spent receiving data during the last sampling period.
- Number of active clients in the last sampling period.
- Activity of the current client in the last sampling period.

The **bw-alloc** parameter of a traffic management profile allows you to set a minimum bandwidth to be allocated to a virtual AP profile when there is congestion on the wireless network. You must set traffic shaping to **fair-access** to use this bandwidth allocation value for an individual virtual AP.

Enabling Traffic Shaping

Traffic shaping is configured in an traffic management profile.

In the WebUI

To configure traffic shaping via the WebUI:

1. Select **Configuration > Advanced Services > All Profiles**. The **All Profile Management** window opens.
2. Select **QoS** to expand the **QoS** section.
3. Select **Traffic management profile**.
4. In the **Profiles Details** window, select the name of the traffic management profile for which you want to configure traffic shaping. (If you do not have any traffic management profiles configured, enter a name for a new profile in the **Profile Details** pane, click **Add**, then select the new profile from the **profiles** list.)
5. In the **Profile Details** pane, click the **Station Shaping Policy** drop-down list and select either **default-access**, **fair-access** or **preferred-access**.
6. Click **Apply**.

The following table describes configuration settings available in the traffic management profile.

Table 83: *Traffic Management Profile Parameters*

Parameter	Description
Station Shaping Policy	Define Station Shaping Policy This feature has the following three options: • **default-access:** Traffic shaping is disabled, and client performance is dependent on MAC contention resolution. This is the default traffic shaping setting. • **fair-access:** Each client gets the same airtime, regardless of client capability and capacity. This option is useful in environments like a training facility or exam hall, where a mix of 802.11a/g, 802.11g, and 802.11n clients need equal to network resources, regardless of their capabilities. The bw-alloc parameter of a traffic management profile allows you to set a minimum bandwidth to be allocated to a virtual AP profile when there is congestion on the wireless network. You must set traffic shaping to fair-access to use this bandwidth allocation value for an individual virtual AP. • **preferred-access:** High-throughput (802.11n) clients do not get penalized because of slower 802.11a/g or 802.11b transmissions that take more air time due to lower rates. Similarly, faster 802.11a/g clients get more access than 802.11b clients.
Proportional BW Allocation	You can allocate a maximum bandwidth, as a percentage of available bandwidth to a virtual AP (VAP). To assign a percentage of bandwidth to a virtual AP: 1. Click the **Virtual AP** drop-down list, and select the VAP to which you would like to allocate a bandwidth share. 2. Specify the percentage of bandwidth to be allocated to the VAP in the **Share(%)** field. 3. Select the **Hard Limit** checkbox to restrict the bandwidth for the VAP. Do not select the **Hard Limit** checkbox if you want to restrict the bandwidth for this VAP when there is a congestion on the wireless network. 4. Click **Add**. 5. Repeat steps 1-4 to assign any remaining bandwidth to additional VAPs, if desired. To remove a VAP from the list of VAPs with allocated bandwidth, select the VAP from the **Proportional BW Allocation** field and click **Delete**.
Report Interval	Number of minutes between bandwidth usage reports. Range: 1-99 minutes Default value is 5 minutes.

In the CLI

To enable and configure traffic shaping via the command-line interface, access the CLI in config mode and issue the following commands:

```
wlan traffic-management-profile <profile> shaping-policy fair-access|preferred-
```

To disable traffic shaping, use the **default-access** parameter:

```
wlan traffic-management-profile <profile> shaping-policy default-access
```

Use the following commands to apply an 802.11a or 802.11g traffic management profile to an AP group or an individual AP.

```
ap-group <name> dot11a-traffic-mgmt-profile|dot11g-traffic-mgmt-profile <profile>
ap-name <name> dot11a-traffic-mgmt-profile|dot11g-traffic-mgmt-profile <profile>
```

Enabling or Disabling the Hard Limit Parameter in Traffic Management Profile

You can configure the limit on OTA bandwidth for a virtual AP by enabling or disabling the hard-limit parameter in the Traffic management profile.

Using the WebUI

The following procedure configures the Hard Limit parameter in Traffic management profile:

1. Navigate to the **Configuration > Advanced Services > All Profiles** page.

2. Under **QOS > Traffic management** on the **Profiles** pane, select the profile name.
3. Under the **Advanced** tab on the **Profile Details** pane, select the **Proportional BW Allocation** parameter and follow the steps given in the Table 83.
4. Click **Apply**.

Using the CLI

You can view the Traffic management profile using the following command:

```
(host)(config) #wlan traffic-management-profile default
```

The following example sets a hard bandwidth limit of 15% for the default virtual-ap under the traffic management profile:

```
(host)(Traffic management profile "default") #bw-alloc virtual-ap default share 15
enforcement hard
```

Spectrum Load Balancing

The spectrum load balancing feature helps optimize network resources by balancing clients across channels, regardless of whether the AP or the controller is responding to the wireless clients' probe requests. The controller uses the ARM neighbor update messages that pass between APs and the controller to determine the distribution of clients connected to each AP's immediate (one-hop) neighbors. This feature also takes into account the number of APs visible to the clients in the RF neighborhood, and can factor the client's perspective on the network into its coverage calculations.

The controller compares whether an AP has more clients than its neighboring APs on other channels. If an AP's client load is at or over a predetermined threshold as compared to its immediate neighbors, or if a neighboring Dell AP on another channel does not have any clients, load balancing will be enabled on that AP.

When an AP has the spectrum load balancing feature enabled, the AP will send an association response with error code 17 to new clients trying to associate. If the client receiving the error code tries to associate to the AP a second time, it will be admitted. If a client is rejected by two APs in a row, it will be admitted by any AP on its third try. Note that the load balancing feature only affects the association of new clients; this feature does not reject or attempt to balance clients that are already associated to the AP.

Spectrum load balancing is disabled by default, and can be enabled for 2.4G traffic through an 802.11g profile or for 5G traffic through an 802.11a RF management profile. The spectrum load balancing feature also requires that the 802.11a or 802.11g RF management profiles reference an ARM profile with ARM scanning enabled.



The spectrum load balancing feature available in ArubaOS 3.4.x and later releases completely replaces the AP load balancing feature available earlier versions of ArubaOS. When you upgrade from an older release to ArubaOS 3.4.x or later, you must manually configure the spectrum load balancing settings, as you can lo longer use the AP load balancing feature, and any previous AP load balancing settings will not be preserved.

For details on modifying 802.11a or 802.11g RF management profiles, refer to RF Management on page 509 .

Reusing Channels to Control RX Sensitivity Tuning

In some dense deployments, it is possible for APs to hear other APs on the same channel. This creates co-channel interference and reduces the overall usage of the channel in a given area. Channel reuse enables dynamic control over the receive (Rx) sensitivity to improve spatial reuse of the channel.



The channel reuse feature applies to non-DFS channels only. It is internally disabled for DFS channels and is does not affect DFS radar signature detection.

You can configure the channel reuse feature to operate in either of the following three modes; *static, dynamic or disable*. (This feature is disabled by default.)

- **Static mode:** This mode of operation is a coverage-based adaptation of the Clear Channel Assessment (CCA) thresholds. In the static mode of operation, the CCA is adjusted according to the configured transmission power level on the AP, so as the AP transmit power decreases as the CCA threshold increases, and vice versa.
- **Dynamic mode:** In this mode, the Clear Channel Assessment (CCA) thresholds are based on channel loads, and take into account the location of the associated clients. When you set the Channel Reuse feature to dynamic mode, this feature is automatically enabled when the wireless medium around the AP is busy greater than half the time, and the CCA threshold adjusts to accommodate transmissions between the AP its most distant associated client.
- **Disable mode:** This mode does not support the tuning of the CCA Detect Threshold.

The channel reuse mode is configured through an 802.11a or 802.11g RF management profile. For details on modifying 802.11a or 802.11g RF management profiles, refer to RF Management on page 509 .

Configuring Non-802.11 Noise Interference Immunity

When an AP attempts to decode a non-802.11 signal, that attempt can momentarily interrupt its ability to receive traffic. The noise immunity feature can help improve network performance in environments with a high level of non-802.11 noise from devices such as Bluetooth headsets, video monitors and cordless phones.

You can configure the noise immunity feature for any one of the following levels of noise sensitivity. Note that increasing the level makes the AP slightly "deaf" to its surroundings, causing the AP to lose a small amount of range.

- Level 0: no ANI adaptation.
- Level 1: Noise immunity only. This level enables power-based packet detection by controlling the amount of power increase that makes a radio aware that it has received a packet.
- Level 2: Noise and spur immunity. This level also controls the detection of OFDM packets, and is the default setting for the Noise Immunity feature.
- Level 3: Level 2 settings and weak OFDM immunity. This level minimizes false detects on the radio due to interference, but may also reduce radio sensitivity. This level is recommended for environments with a high-level of interference related to 2.4Ghz, appliances such as cordless phones.
- Level 4: Level 3 settings, and FIR immunity. At this level, the AP adjusts its sensitivity to in-band power, which can improve performance in environments with high, constant levels of noise interference.
- Level 5: The AP completely disables PHY error reporting, improving performance by eliminating the time the controller would spend on PHY processing.



Only 802.11n-capable APs simultaneously support both the RX Sensitivity Tuning Based Channel Reuse feature and a level-3 to level-5 Noise Immunity setting. Do not raise the noise immunity default setting on APs that do not support 802.11n unless you first disable the Channel Reuse feature.

You can manage Non-802.11 Noise Immunity settings through the **Non 802.11 Interference Immunity** parameter in the 802.11a or 802.11g RF management profile. For details on configuring this profile, refer to RF Management on page 509

Troubleshooting ARM

If the APs on your WLAN do not seem to be operating at an optimal channel or power setting, you should first verify that both the ARM feature and ARM scanning have been enabled. Optimal ARM performance requires

that the APs have IP connectivity to their master controller, as it is the master controller that gives each AP the global classification information required to keep accurate coverage index values. If ARM is enabled but does not seem to be working properly, try some of the troubleshooting tips below.

Too many APs on the Same Channel

If many APs are selecting the same RF channel, there may be excessive interference on the other valid 802.11 channels. Issue the CLI commands

`show ap arm rf-summary ap-name <ap-name> or show ap arm rf-summary ip-addr <ap ip address>` and calculate the Interference index (*intf_idx*) for all the valid channels.

An AP will only move to a new channel if the new channel has a lower interference index value than the current channel. The ARM Free Channel Index parameter specifies the required difference between two interference index values. If this value is set too high, the AP will not switch channels, even if the interference is slightly lower on another channel. Lower the Free Channel Index to improve the likelihood that the AP will switch to a better channel.

Wireless Clients Report a Low Signal Level

If APs detect strong signals from other APs on the same channel, they may decrease their power levels accordingly. Issue the CLI commands

`show ap arm rf-summary ap-name <ap-name>` or **`show ap arm rf-summary ip-addr <ap ip address>`**

for all APs and check their current coverage index (*cov-idx*). If the AP's coverage index is at or higher than the configured coverage index value, then the APs have correctly chosen the transmit power setting. To manually increase the minimum power level for the APs using a specific ARM profile, define a higher minimum value with the following command:

`rf arm-profile <profile> min-tx-power <dBm>`.

If wireless clients still report that they see low signal levels for the APs, check that the AP's antennas are correctly connected to the AP and correctly placed according to the manufacturer's installation guide.

Transmission Power Levels Change Too Often

Frequent changes in transmission power levels can indicate an unstable RF environment, but can also reflect incorrect ARM or AP settings. To slow down the frequency at which the APs change their transmit power, set the ARM Backoff Time to a higher value. If APs use external antennas, check the **Configuration > Wireless > AP Installation > Provisioning** window to make sure the APs are statically configured for the correct dBi gain, antenna type, and antenna number. If only one external antenna is connected to its radio, you must select either antenna number 1 or 2.

APs Detect Errors but Do Not Change Channels

First, ensure that ARM error checking is enabled. The ARM Error Rate Threshold should be set to a percentage higher than zero. The suggested configuration value for the ARM Error Rate Threshold is 30–50%.

APs Don't Change Channels Due to Channel Noise

APs will only change channels due to interference if you enable ARM noise checking. Check to verify that the ARM Noise Threshold is set to a value higher than 0 dBm. The suggested setting for this threshold is 75 dBm.

The ArubaOS Wireless Intrusion Prevention (WIP) features and configurations are discussed in this chapter. WIP offers a wide selection of intrusion detection and protection features to protect the network against wireless threats. Like most other security-related features of the Dell network, the WIP configuration is done on the master controller in the network.

To use most of the features described in this chapter, you must install a Wireless Intrusion Protection (RFprotect) license on all controllers in your network. If you install a RFprotect license on a master controller only, an AP or AM terminated on a local controller will not provide the WIP features.

These features do not require an RFprotect license:

- Rogue AP classification techniques other than AP classification rules
- Rogue containment
- Wired containment
- Wireless containment without Tarpit

For details on commands see the *Dell Networking W-Series ArubaOS 6.4.x Command Line Interface Guide*.

This chapter contains the following sections:

- Working with the Reusable Wizard on page 452
- Monitoring the Dashboard on page 455
- Detecting Rogue APs on page 456
- Working with Intrusion Detection on page 459
- Configuring Intrusion Protection on page 471
- Configuring the WLAN Management System (WMS) on page 475
- Understanding Client Blacklisting on page 476
- Working with WIP Advanced Features on page 479
- Configuring TotalWatch on page 480
- Administering TotalWatch on page 482
- Tarpit Shielding Overview on page 483
- Configuring Tarpit Shielding on page 484

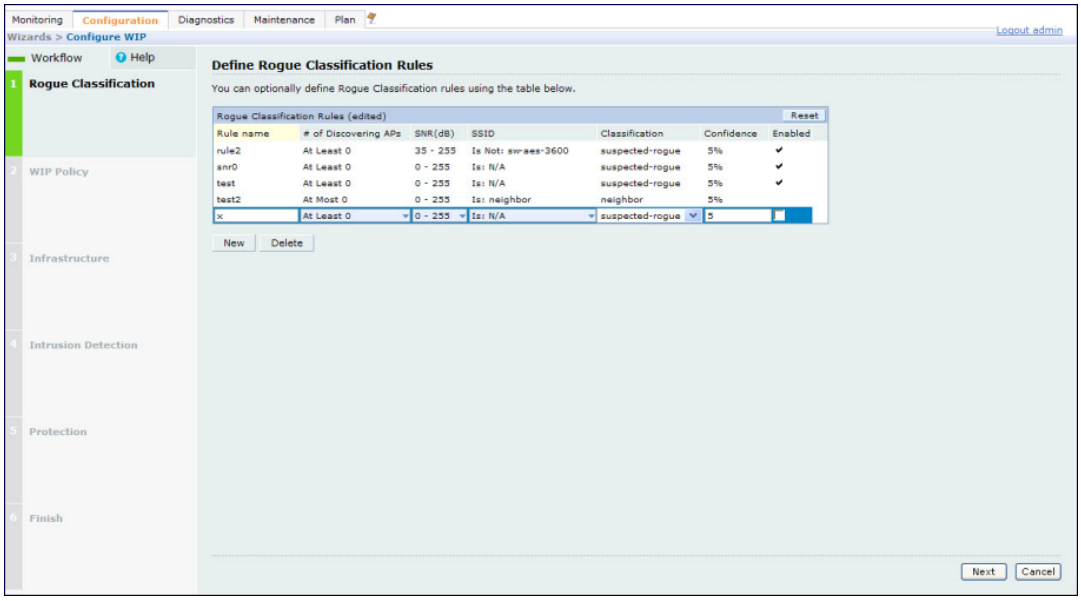
Working with the Reusable Wizard

The WebUI's reusable, intuitive, user-friendly Wizard provides steps to enable, define, or change:

- Integrated vs Overlay WLAN/WIP options
- Rules-based rogue classification
- Detection features for attacks against infrastructure
- Detection features for attacks against WLAN clients
- Protection features for attacks against infrastructure
- Protection features for WLAN clients

Figure 46 displays the WIP Wizard layout. Highlighting one of the previously configured rules reveals drop down menus for changing values. Note that the reusable wizard includes robust online Help.

Figure 46 *WIP Wizard*



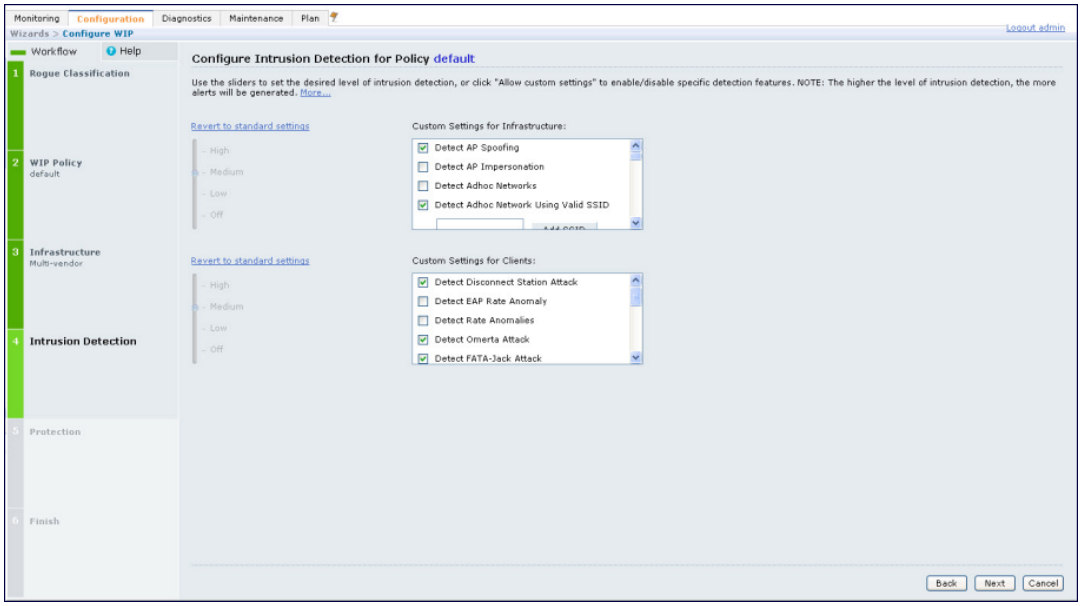
Understanding Wizard Intrusion Detection

Apply the intrusion detection mechanisms for detecting attacks against your infrastructure and clients (see Figure 47). You can either set the detection level to automatically enable the appropriate detection mechanisms or customize the settings for infrastructure and client attacks. Use the slider to select one of the detection levels for the infrastructure and clients:

- High—Enables all the detection mechanisms applicable to your infrastructure including all the options of low and medium level settings.
- Medium (Default)—Enables some important detection mechanisms for your infrastructure. This includes all the options of the low level settings.
- Low—Enables only the most critical detection mechanisms for your infrastructure.
- Off—Disables all the detection mechanisms.

To enable custom settings, click the *Allow custom settings* link to manually enable or disable the detection mechanisms for your clients. To revert to the standard settings from the custom settings mode, click the *Revert to standard settings* link.

Figure 47 *WIP Wizard's Intrusion Detection*



Understanding Wizard Intrusion Protection

Apply the intrusion protection mechanisms for your infrastructure and clients (see Figure 48). You can set the protection level to automatically enable the appropriate protection mechanisms or customize the settings for your infrastructure and clients.

Protecting Your Infrastructure

Use the slider to select one of the protection levels for the infrastructure:

- High—Enables all the protection mechanisms applicable to your infrastructure including all the options of low and medium level settings.
- Medium—Enables some important protection mechanisms for your infrastructure including all the options of the low level settings.
- Low—Enables only the most critical protection mechanisms for your infrastructure.
- Off (Default)—Disables all the protection mechanisms.

To enable custom settings, click the *Allow custom settings* link. You can manually enable or disable the protection mechanisms for your infrastructure. To revert to the standard settings from custom settings mode, click the *Revert to standard settings* link.

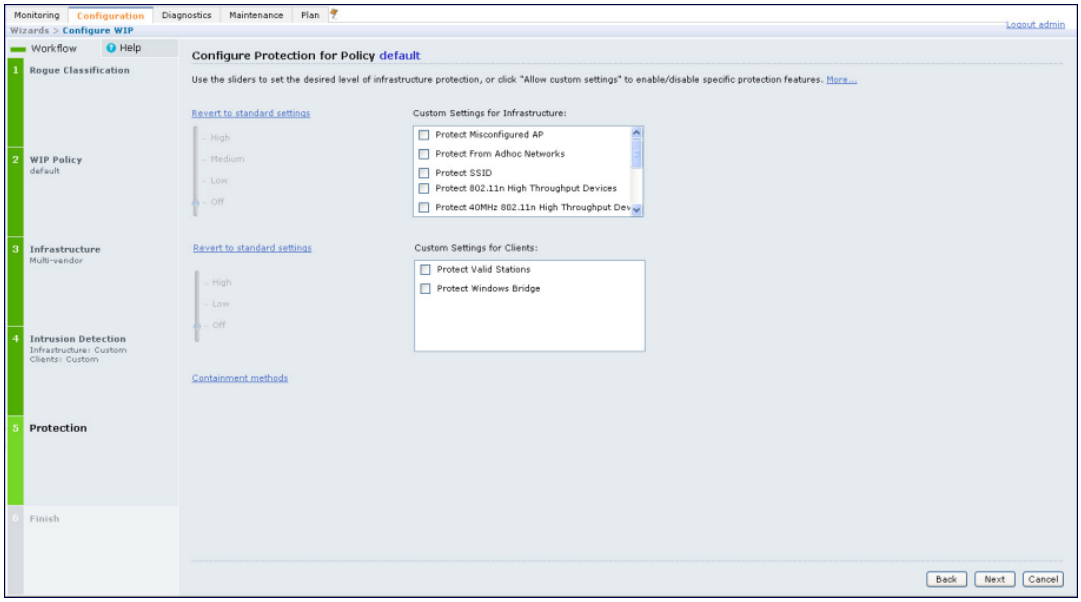
Protecting Your Clients

Use the slider (see Figure 48) to select one of the following preset protection levels for your clients:

- High—Enables all the protection mechanisms applicable to your clients including all the options of the low level settings.
- Low—Enables only the most critical protection mechanisms for your clients.
- Off (Default)—Disables all the protection mechanisms.

To enable custom settings, click the *Allow custom settings* link to manually enable or disable the protection mechanisms for your clients. To revert to the standard settings from custom settings mode, click the *Revert to standard settings* link.

Figure 48 *WIP Wizard Intrusion Protection*



Monitoring the Dashboard

The Security Summary dashboard, in the Monitoring section of the WebUI, allows you to monitor the detection and protection of wireless intrusions in your network.

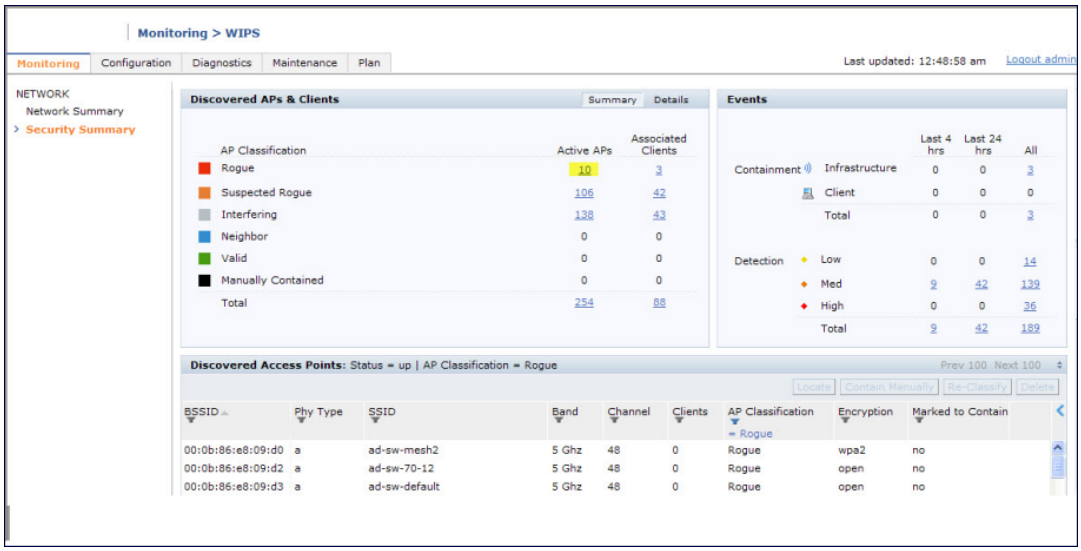
The dashboard's two top tables—Discovered APs & Clients and Events—contain data as links. When these links are selected they arrange, filter, and display the appropriate information in the lower table. For example, if you select the number 10 under the Active APs column (highlighted in yellow in Figure 49) then the bottom table will filter and arrange information about the ten classified Rogue APs. Use the scroll bar at the right to view all ten Rogue APs.



The term *events* in this document is meant to include security threats, vulnerabilities, attacks (intrusion or Denial of Service) and other similarly related events.

The Event table contains data links. Selecting these data links will display information, in the bottom table, related to the Event you selected. Again, remember to use the scroll bar at the right to view all the Events.

Figure 49 *WIP Monitoring Dashboard*



Detecting Rogue APs

The most important WIP functionality is the ability to classify an AP as a potential security threat. An AP is considered to be a rogue AP if it is both unauthorized and plugged into the wired side of the network. An AP is considered to be an interfering AP if it is seen in the RF environment but is not connected to the wired network.

While the interfering AP can potentially cause RF interference, it is not considered a direct security threat since it is not connected to the wired network. However, an interfering AP may be reclassified as a rogue AP.

Understanding Classification Terminology

APs and clients are discovered during scanning of the wireless medium, and they are classified into various groups. The AP and client classification definitions are in Table 84 and Table 85.

Table 84: *AP Classification Definition*

Classification	Description
Valid AP	An AP that is part of the enterprise providing WLAN service.
Interfering AP	An AP that is seen in the RF environment but is not connected to the wired network. An interfering AP is not considered a direct security threat since it is not connected to the wired network. For example, an interfering AP can be an AP that belongs to a neighboring office's WLAN but is not part of your WLAN network.
Neighbor AP	A neighboring AP is when the BSSIDs are known. Once classified, a neighboring AP does not change its state.
Rogue AP	An unauthorized AP that is plugged into the wired side of the network.
Suspected-Rogue AP	A suspected rogue AP is an unauthorized AP that may be plugged into the wired side of the network.
Manually-contained AP	An AP for which DoS is enabled manually.

Table 85: *Client Classification Definitions*

Classification	Description
Valid Client	Any client that successfully authenticates with a valid AP and passes encrypted traffic is classified as a valid client.
Manually-contained Client	Any clients for which DoS is enabled manually.
Interfering Client	A client associated to any AP and is not valid.

Understanding Classification Methodology

A discovered AP is classified as a rogue or a suspected rogue by the following methods:

- Internal heuristics
- AP classification rules
- Manually by the user

The internal heuristics works by checking if the discovered AP is communicating with a wired device on the customer network. This is done by matching the MAC address of devices that are on the discovered AP's network with that of the user's wired network. The MAC of the device on the discovered AP's network is known as the *Match MAC*. The ways in which the matching of wired MACs occurs is detailed in the sections Understanding Match Methods on page 457 and Understanding Match Types on page 457.

Understanding Match Methods

The match methods are:

- Plus One—The match MAC matches a device whose MAC address' last bit was one more than that of the Match MAC.
- Minus One—The match MAC matches a device whose MAC address' last bit was one less than that of the Match MAC.
- Equal—The match was against the same MAC address.
- OUI—The match was against the manufacturer's OUI of the wired device.

The classification details are available in the 'Discovered AP table' section of the 'Security Summary' page of the WebUI. The information can be obtained by clicking on the details icon for a selected discovered AP. The information is also available in the command **show wms rogue-ap**.

Understanding Match Types

- Eth-Wired-MAC—The MAC addresses of wired devices learned by an AP on its Ethernet interface.
- GW-Wired-MAC—The collection of Gateway MACs of all APs across the master and local controllers.
- AP-Wired-MAC—The MAC addresses of wired devices learned by monitoring traffic out of other valid and rogue APs.
- Config-Wired-MAC—The MAC addresses that are configured by the user typically that of well known servers in the network.
- Manual—User triggered classification.
- External-Wired-MAC—The MAC address matched a set of known wired devices that are maintained in an external database.
- Mobility-Manager—The classification was determined by the mobility manager, AMP.

- Classification-off—AP is classified as rogue because classification has been disabled causing all non-authorized APs to be classified as a rogue.
- Propagated-Wired-MAC—The MAC addresses of wired devices learned by a different AP than the one that uses it for classifying a rogue.
- Base-BSSID-Override—The classification was derived from another BSSID which belongs to the same AP that supports multiple BSSIDs on the radio interface.
- AP-Rule—A user defined AP classification rule has matched.
- System-Wired-MAC—The MAC addresses of wired devices learned at the controller.
- System-Gateway-MAC—The Gateway MAC addresses learned at the controller.

Understanding Suspected Rogue Confidence Level

A suspected rogue AP is an AP that is potentially a threat to the WLAN infrastructure. A suspected rogue AP has a confidence level associated with it. An AP can be marked as a suspected rogue if it is determined to be a potentially threat on the wired network, or if it matches a user defined classification rule.

The suspected-rogue classification mechanism are:

- Each mechanism that causes a suspected-rogue classification is assigned a confidence level increment of 20%.
- AP classification rules have a configured confidence level.
- When a mechanism matches a previously unmatched mechanism, the confidence level increment associated with that mechanism is added to the current confidence level (the confident level starts at zero).
- The confidence level is capped at 100%.
- If your controller reboots, your suspected-rogue APs are not checked against any new rules that were configured after the reboot. Without this restriction, all the mechanisms that classified your APs as suspected-rogue may trigger again causing the confidence level to surpass their cap of 100%. You can explicitly mark an AP as “interfering” to trigger all new rules to match against it.

Understanding AP Classification Rules

AP classification rule configuration is performed only on a master controller. If AMP is enabled via the mobility-manager command, then processing of the AP classification rules is disabled on the master controller. A rule is identified by its ASCII character string name (32 characters maximum). The AP classification rules have one of the following specifications:

- SSID of the AP
- SNR of the AP
- Discovered-AP-Count or the number of APs that can see the AP

Understanding SSID specification

Each rule can have up to 6 SSID parameters. If one or more SSIDs are specified in a rule, an option of whether to match any of the SSIDs, or to not match all of the SSIDs can be specified. The default is to check for a match operation.

Understanding SNR specification

Each rule can have only one specification of the SNR. A minimum and/or maximum can be specified in each rule and the specification is in SNR (db).

Understanding Discovered-AP-Count specification

Each rule can have only one specification of the Discovered-AP-Count. Each rule can specify a minimum or maximum of the Discovered-AP-count. The minimum or maximum operation must be specified if the Discovered-AP-count is specified. The default setting is to check for the minimum discovered-AP-count.

Sample Rules

If SSID equals xyz AND SNR > 40 then classify AP as suspected-rogue with conf-level-increment of 20

If SNR > 60 and DISCOVERING_APS > 2, then classify AP as suspected-rogue with conf-level increment of 35

If SSID equals 'XYZ', then classify AP as known-neighbor

Understanding Rule Matching

A rule must be enabled before it is matched. A maximum of 32 rules can be created with a maximum of 16 rules active simultaneously. If a rule matches, an AP is classified to:

- Suspected-Rogue—an associated confidence-level is provided (minimum is 5%)
- Neighbor

The following mechanism is used for rule matching.

- When *all* the conditions specified in the rule evaluate to true, the rule matches.
- If multiple rules match causing the AP to be classified as a Suspected-Rogue, the confidence level of each rule is aggregated to determine the confidence level of the classification.
- When multiple rules match and any one of those matching rules cause the AP to be classified as a Neighbor, then the AP is classified as Neighbor.
- APs classified as either Neighbor or Suspected-Rogue will attempted to match any configured AP rule.
- Once a rule matches an AP, the same rule will not be checked for the AP.
- When the controller reboots, no attempt to match a previously matched AP is made.
- If a rule is disabled or modified, all APs that were previously classified based on that rule will continue to be in the newly classified state.

Working with Intrusion Detection

This section covers Infrastructure and Client Intrusion Detections.

Understanding Infrastructure Intrusion Detection

Detecting attacks against the infrastructure is critical in avoiding attacks that may lead to a large-scale Denial of Service (DOS) attack or a security breach. This group of features detects attacks against the WLAN infrastructure, which consists of authorized APs, the RF medium, and the wired network. An authorized or valid-AP is defined as an AP that belongs to the WLAN infrastructure. The AP is either a Dell AP or a third party AP. ArubaOS automatically learns authorized Dell APs.

Table 86 presents a summary of the Intrusion infrastructure detection features with their related commands, traps, and syslog identification. Feature details follow the table.

Table 86: *Infrastructure Detection Summary*

Feature	Command	Trap	Syslog ID
Detecting an 802.11n 40MHz Intolerance Setting on page 463	`ids dos-profile detect-ht-40mhz-intolerance client-ht-40mhz-intol-quiet-time`	wlsxHT40MHzIntoleranceAP wlsxHT40MHzIntoleranceSta	126052, 126053, 127052, 127053
Detecting Active 802.11n Greenfield Mode on page 464	`ids unauthorized-device-profile detect-ht-greenfield`	wlsxHtGreenfieldSupported	126054, 127054
Detecting Ad hoc Networks on page 464	`ids unauthorized-device-profile detect-adhoc-network`	wlsxNAdhocNetwork	126033, 127033
Detecting an Ad hoc Network Using a Valid SSID on page 464	`ids unauthorized-device-profile detect-adhoc-using-valid-ssid adhoc-using-valid-ssid-quiet-time`	wlsxAdhocUsingValidSSID	126068, 127068
Detecting an AP Flood Attack on page 464	`ids dos-profile detect-ap-flood ap-flood-threshold ap-flood-inc-time ap-flood-quiet-time`	wlsxApFloodAttack	126034, 127034
Detecting AP Impersonation on page 464	`ids impersonation-profile detect-ap-impersonation beacon-diff-threshold beacon-inc-wait-time`	wlsxAPImpersonation	126006, 127006
Detecting AP Spoofing on page 464	`ids impersonation-profile detect-ap-spoofing ap-spoofing-quiet-time`	wlsxAPSpoofingDetected wlsxClientAssociatingOnWrongChannel	126069, 126070, 127069, 127070

Feature	Command	Trap	Syslog ID
Detecting Bad WEP Initialization on page 464	`ids unauthorized-device-profile detect-bad-wep`	wlsxRepeatWEPIVViolation wlsxStaRepeatWEPIVViolation wlsxWeakWEPIVViolation wlsxStaWeakWEPIVViolation	126014, 126015, 126016, 126017, 127014, 127015, 127016, 127017
Detecting a Beacon Frame Spoofing Attack on page 464	`ids impersonation-profile detect-beacon-wrong-channel beacon-wrong-channel-quiet-time`	wlsxMalformedFrameWrongChannel Detected	126086, 127086
Detecting a Client Flood Attack on page 464	`ids dos-profile detect-client-flood client-flood-threshold client-flood-inc-time client-flood-quiet-time`	wlsxClientFloodAttack	126064, 127064
Detecting a CTS Rate Anomaly	`ids dos-profile detect-cts-rate-anomaly cts-rate-threshold cts-rate-time-interval cts-rate-quiet-time`	wlsxCtsRateAnomaly	126073, 127073
Detecting Devices with an Invalid MAC OUI on page 465	`ids unauthorized-device-profile detect-invalid-mac-oui mac-oui-quiet-time`	wlsxInvalidMacOUIAP wlsxInvalidMacOUISta	126029, 126030, 127029, 127030
Detecting an Invalid Address Combination on page 465	`ids dos-profile detect-invalid-address-combination invalid-address-combination-quiet-time`	wlsxInvalidAddressCombination	126079, 127079
Detecting an Overflow EAPOL Key on page 465	`ids dos-profile detect-overflow-eapol-key overflow-eapol-key-quiet-time`	wlsxMalformedOverflowEAPOLKey Detected	126082, 127082
Detecting Overflow IE Tags on page 465	`ids dos-profile detect-overflow-ie overflow-ie-quiet-time`	wlsxOverflowIEDetected	126084, 127084

Feature	Command	Trap	Syslog ID

Detecting a Malformed Frame-Assoc Request on page 465	`ids dos-profile`
	`detect-malformed-assoc-req`
	`malformed-assoc-req-quiet-time`
		wlsxMalformedAssocReqDetected	126080, 127080

Detecting Malformed Frame-Auth on page 465	`ids dos-profile`
	`detect-malformed-frame-auth`
	`malformed-auth-frame-quiet-time`
		wlsxMalformedAuthFrameDetected	126083, 127083

Detecting a Malformed Frame-HT IE on page 466	`ids dos-profile`
	`detect-malformed-htie`
	`malformed-htie-quiet-time`
		wlsxMalformedHTIEDetected	126081, 127081

Detecting a Malformed Frame-Large Duration on page 466	`ids-dos-profile`
	`detect-malformed-large-duration`
	`malformed-large-duration-quiet-time`
		wlsxMalformedFrameLargeDuration Detected	126085, 127085

Detecting a Misconfigured AP on page 466 (WEP, WPA, SSID, Channel, OUI)	`ids unauthorized-device-profile`
	`detect-misconfigured-ap`
	`privacy`
	`require-wpa`
	`valid-and-protected-ssid`
	`cfg-valid-11g-channel`
	`cfg-valid-11a-channel`
	`valid-oui`
		wlsxWEPMisconfiguration
		wlsxWPAMisconfiguration
		wlsxSSIDMisconfiguration
		wlsxChannelMisconfiguration
		wlsxOUIMisconfiguration
			126011, 126028, 126010, 126008, 126009, 127011, 127028, 127010, 127008, 127009

Detecting a CTS Rate Anomaly on page 465	`ids dos-profile`
	`detect-rts-rate-anomaly`
	`rts-rate-threshold`
	`rts-rate-time-interval`
	`rts-rate-quiet-time`
		wlsxRtsRateAnomaly	126074, 127074

Feature	Command	Trap	Syslog ID
Detecting a Windows Bridge on page 466	`ids unauthorized-device-profile detect-windows-bridge`	wlsxWindowsBridgeDetectedAP wlsxWindowsBridgeDetectedSta wlsxNAdhocNetworkBridgeDetected AP wlsxNAdhocNetworkBridgeDetected Sta	126039, 126040, 126041, 126042, 127039, 127040, 127041, 127042
Detecting a Wireless Bridge on page 466	`ids unauthorized-device-profile detect-wireless-bridge wireless-bridge-quiet-time`	wlsxWirelessBridge	126036, 127036

Detecting Broadcast Deauthentication on page 466	`ids signature-matching-profile signature deauth-Broadcast`
	
	`ids general-profile signature-quiet-time`
		wlsxNSignatureMatchDeauthBcast	126047, 127047

Detecting Broadcast Disassociation on page 466	`ids signature-matching-profile signature disassoc-Broadcast`
	
	`ids general-profile signature-quiet-time`
		wlsxNSignatureMatchDisassocBcast	126066, 127066

Detecting Netstumbler on page 466	`ids signature-matching-profile signature 'Netstumbler Generic' signature 'Netstumbler Version 3.3.0.x'`
	
	`ids general-profile signature-quiet-time`
		wlsxNSignatureMatchNetstumbler	126043, 127043

Detecting Valid SSID Misuse on page 466	`ids-unauthorized-device-profile detect-valid-ssid-misuse valid-and-protected-ssid`	wlsxValidSSIDViolation	126007, 127007

Detecting Wellenreiter on page 466	`ids signature-matching-profile signature Wellenreiter`
	
	`ids general-profile signature-quiet-time`
		wlsxNSignatureMatchWellenreiter	126067, 127067

Detecting an 802.11n 40MHz Intolerance Setting

When a client sets the HT capability "**intolerant** bit" to indicate that it is unable to participate in a 40MHz BSS, the AP must use lower data rates with all of its clients. Network administrators often want to know if there are devices that are advertising 40MHz intolerance, as this can impact the performance of the network.

Detecting Active 802.11n Greenfield Mode

When 802.11 devices use the HT operating mode, they can not share the same channel as 802.11a/b/g stations. Not only can they not communicate with legacy devices, the way they use the transmission medium is different, which would cause collisions, errors and retransmissions.

Detecting Ad hoc Networks

An ad hoc network is a collection of wireless clients that form a network amongst themselves without the use of an AP. As far as network administrators are concerned, ad hoc wireless networks are uncontrolled. If they do not use encryption, they may expose sensitive data to outside eavesdroppers. If a device is connected to a wired network and has bridging enabled, an ad-hoc network may also function like a rogue AP. Additionally, ad-hoc networks can expose client devices to viruses and other security vulnerabilities. For these reasons, many administrators choose to prohibit ad-hoc networks.

Detecting an Ad hoc Network Using a Valid SSID

If an unauthorized ad hoc network is using the same SSID as an authorized network, a valid client may be tricked into connecting to the wrong network. If a client connects to a malicious ad hoc network, security breaches or attacks can occur.

Detecting an AP Flood Attack

Fake AP is a tool that was originally created to thwart wardrivers by flooding beacon frames containing hundreds of different addresses. This would appear to a wardriver as though there were hundreds of APs in the area, thus concealing the real AP. An attacker can use this tool to flood an enterprise or public hotspots with fake AP beacons to confuse legitimate users and to increase the amount of processing need on client operating systems.

Detecting AP Impersonation

In AP impersonation attacks, the attacker sets up an AP that assumes the BSSID and ESSID of a valid AP. AP impersonation attacks can be done for man-in-the-middle attacks, a rogue AP attempting to bypass detection, or a honeypot attack.

Detecting AP Spoofing

An AP Spoofing attack involves an intruder sending forged frames that are made to look like they are from a legitimate AP. It is trivial for an attacker to do this, since tools are readily available to inject wireless frames with any MAC address that the user desires. Spoofing frames from a legitimate AP is the foundation of many wireless attacks.

Detecting Bad WEP Initialization

This is the detection of WEP initialization vectors that are known to be weak. A primary means of cracking WEP keys is to capture 802.11 frames over an extended period of time and searching for such weak implementations that are still used by many legacy devices.

Detecting a Beacon Frame Spoofing Attack

In this type of attack, an intruder spoofs a beacon packet on a channel that is different from that advertised in the beacon frame of the AP.

Detecting a Client Flood Attack

There are fake AP tools that can be used to attack wireless intrusion detection itself by generating a large number of fake clients that fill internal tables with fake information. If successful, it overwhelms the wireless intrusion system, resulting in a DoS.

Detecting a CTS Rate Anomaly

The RF medium can be reserved via Virtual Carrier Sensing using an Clear To Send (CTS) transaction. The transmitter station sends a Ready To Send (RTS) frame to the receiver station. The receiver station responds with a CTS frame. All other stations that receive these CTS frames will refrain from transmitting over the wireless medium for an amount of time specified in the *duration* fields of these frames.

Attackers can exploit the Virtual Carrier Sensing mechanism to launch a DoS attack on the WLAN by transmitting numerous RTS and/or CTS frames. This causes other stations in the WLAN to defer transmission to the wireless medium. The attacker can essentially block the authorized stations in the WLAN with this attack.

Detecting an RTS Rate Anomaly

The RF medium can be reserved via Virtual Carrier Sensing using an RTS transaction. The transmitter station sends a RTS frame to the receiver station. The receiver station responds with a CTS frame. All other stations that receive these RTS frames will refrain from transmitting over the wireless medium for an amount of time specified in the *duration* fields of these frames.

Attackers can exploit the Virtual Carrier Sensing mechanism to launch a DoS attack on the WLAN by transmitting numerous RTS and/or CTS frames. This causes other stations in the WLAN to defer transmission to the wireless medium. The attacker can essentially block the authorized stations in the WLAN with this attack.

Detecting Devices with an Invalid MAC OUI

The first three bytes of a MAC address, known as the MAC organizationally unique identifier (OUI), is assigned by the IEEE to known manufacturers. Often clients using a spoofed MAC address do not use a valid OUI and instead use a randomly generated MAC address.

Detecting an Invalid Address Combination

In this attack, an intruder can cause an AP to transmit deauthentication and disassociation frames to all of its clients. Triggers that can cause this condition include the use of broadcast or multicast MAC address in the source address field.

Detecting an Overflow EAPOL Key

Some wireless drivers used in access points do not correctly validate the EAPOL key fields. A malicious EAPOL-Key packet with an invalid advertised length can trigger a DoS or possible code execution. This can only be achieved after a successful 802.11 association exchange.

Detecting Overflow IE Tags

Some wireless drivers used in access points do not correctly parse the vendor-specific IE tags. A malicious association request sent to the AP containing an IE with an inappropriate length (too long) can cause a DoS and potentially lead to code execution. The association request must be sent after a successful 802.11 authentication exchange.

Detecting a Malformed Frame-Assoc Request

Some wireless drivers used in access points do not correctly parse the SSID information element tag contained in association request frames. A malicious association request with a null SSID (that is, zero length SSID) can trigger a DoS or potential code execution condition on the targeted device.

Detecting Malformed Frame-Auth

Malformed 802.11 authentication frames that do not conform to the specification can expose vulnerabilities in some drivers that have not implemented proper error checking. This feature checks for unexpected values in a Authentication frame.

Detecting a Malformed Frame-HT IE

The IEEE 802.11n HT (High Throughput) IE is used to convey information about the 802.11n network. A 802.11 management frame containing a malformed HT IE can crash some client implementations; potentially representing an exploitable condition when transmitted by a malicious attacker.

Detecting a Malformed Frame-Large Duration

The virtual carrier-sense attack is implemented by modifying the 802.11 MAC layer implementation to allow random duration values to be sent periodically. This attack can be carried out on the ACK, data, RTS, and CTS frame types by using large duration values. This attack can prevent channel access to legitimate users.

Detecting a Misconfigured AP

A list of parameters can be configured that defines the characteristics of a valid AP. This feature is primarily used when non-Dell APs are used in the network since the Dell controller cannot configure the third-party APs. These parameters include WEP, WPA, OUI of valid MAC addresses, valid channels, and valid SSIDs.

Detecting a Windows Bridge

A Windows Bridge occurs when a client that is associated to an AP is also connected to the wired network, and has enabled bridging between these two interfaces.

Detecting a Wireless Bridge

Wireless bridges are normally used to connect multiple buildings together. However, an attacker could place (or have an authorized person place) a wireless bridge inside the network that would extend the corporate network somewhere outside the building. Wireless bridges are somewhat different from rogue APs in that they do not use beacons and have no concept of association. Most networks do not use bridges – in these networks, the presence of a bridge is a signal that a security problem exists.

Detecting Broadcast Deauthentication

A deauthentication broadcast attempts to disconnect all stations in range. Rather than sending a spoofed deauth to a specific MAC address, this attack sends the frame to a broadcast address.

Detecting Broadcast Disassociation

By sending disassociation frames to the broadcast address (FF:FF:FF:FF:FF:FF), an attacker can disconnect all stations on a network for a widespread DoS.

Detecting Netstumbler

NetStumbler is a popular wardriving application used to locate 802.11 networks. When used with certain NICs, NetStumbler generates a characteristic frame that can be detected. Version 3.3.0 of NetStumbler changed the characteristic frame slightly.

Detecting Valid SSID Misuse

If an unauthorized AP (neighbor or interfering) is using the same SSID as an authorized network, a valid client may be tricked into connecting to the wrong network. If a client connects to a malicious network, security breaches or attacks can occur.

Detecting Wellenreiter

Wellenreiter is a passive wireless network discovery tool that is used to compile a list of APs along with their MAC address, SSID, channel, security setting in the vicinity. It passively sniffs wireless traffic and with certain version (versions 1.4, 1.5, and 1.6) sends active probes that target known default SSIDs.

Understanding Client Intrusion Detection

Generally, clients are more vulnerable to attacks than APs. Clients are more apt to associate with a malignant AP due to the client's driver behavior or to a misconfigured client. It is important to monitor authorized clients to track their associations and to track any attacks raised against the client.Client attack detection is categorized as:

- Detecting attacks against Dell APs clients—An attacker can perform an active DOS attack against an associated client, or perform a replay attack to obtain the keys of transmission which could lead to more serious attacks.
- Monitoring Authorized clients—Since clients are easily tricked into associating with unauthorized APs, tracking all misassociations of authorized clients is very important.

An authorized client is a client authorized to use the WLAN network. In ArubaOS, an authorized client is called a *valid-client*. ArubaOS automatically learns a valid client. A client is determined to be valid if it is associated to an authorized or valid AP using encryption; either Layer 2 or IPSEC.



Detection of attacks is limited to valid clients and clients associated to valid APs. Clients that are associated as guests using unencrypted association are included in the attack detection. However, clients on neighboring (interfering) APs are not tracked for attack detection unless they are specified as valid.

Table 87 presents a summary of the client intrusion detection features with their related commands, traps, and syslog identification. Details of each feature follow the table.

Table 87: *Client Detection Summary*

Feature	Command	Trap	Syslog ID
Detecting a Block ACK DoS on page 469	`ids-dos-profile detect-block-ack-attack block-ack-quiet-time`	wlsxBlockAckAttackDetected	126087, 127087
Detecting a ChopChop Attack on page 469	`ids-dos-profile detect-chopchop-attack chopchop-quiet-time`	wlsxChopChopAttackDetected	126078, 127078
Detecting a Disconnect Station Attack on page 469	`ids dos-profile <name> detect-disconnect-sta disconnect-sta-quiet-time disconnect-sta-assoc-resp-threshold disconnect-deauth-disassoc-threshold`	wlsxNDisconnectStationAttack	126035, 127035
Detecting an EAP Rate Anomaly on page 469	`ids-dos-profile detect-eap-rate-anomaly eap-rate-threshold eap-rate-time-interval eap-rate-quiet-time`	wlsxEAPRateAnomaly	126032, 127032
Detecting a FATA-Jack Attack Structure on page 469	`ids dos-profile detect-fatajack-attack fatajack-attack-quiet-time`	wlsxFataJackAttackDetected	126072, 127072

Feature	Command	Trap	Syslog ID
Detecting a Hotspotter Attack on page 470	`ids impersonation-profile detect-hotspotter-attack hotspotter-quiet-time`	wlsxHotspotterAttackDetected	126088, 127088
Detecting a Meiners Power Save DoS Attack on page 470	`ids dos-profile detect-power-save-dos-attack power-save-dos-min-frames power-save-dos-quiet-time power-save-dos-threshold`	wlsxPowerSaveDoSAttack	126109, 127109
Detecting an Omerta Attack on page 470	`ids dos-profile detect-omerta-attack omerta-attack-threshold omerta-attack-quiet-time`	wlsxOmertaAttack	126071, 127071

Detecting Rate Anomalies on page 470	`ids dos-profile detect-rate-anomalies`
	
	`assoc-rate-thresholds disassoc-rate-thresholds deauth-rate-thresholds probe-request-rate-thresholds probe-response-rate-thresholds auth-rate-thresholds`
		wlsxChannelRateAnomaly
		wlsxNodeRateAnomalyAP
		wlsxNodeRateAnomalySta
			126061, 126062, 126063, 127061, 127062, 127063

Detecting a TKIP Replay Attack on page 470	`ids dos-profile detect-tkip-replay-attack tkip-replay-quiet-time`	wlsxTkipReplayAttackDetected	126077, 127077
Detecting Unencrypted Valid Clients on page 470	`ids unauthorized-device-profile detect-unencrypted-valid-client unencrypted-valid-client-quiet-time`	wlsxValidClientNotUsingEncryption	126065, 127065
Detecting a Valid Client Misassociation on page 470	`ids unauthorized-device-profile detect-valid-client-misassociation`	wlsxValidClientMisassociation	126075, 127075

Feature	Command	Trap	Syslog ID

Detecting an AirJack Attack on page 471	`ids signature-matching-profile signature AirJack`
	
	`ids general-profile signature-quiet-time`
		wlsxNSignatureMatchAirjack	126046, 127046

Detecting ASLEAP on page 471	`ids signature-matching-profile signature ASLEAP`
	
	`ids general-profile signature-quiet-time`
		wlsxNSignatureMatchAsleap	126044, 127044

Detecting a Null Probe Response on page 471	`ids signature-matching-profile signature Null Probe Response`
	
	`ids general-profile signature-quiet-time`
		wlsxNSignatureMatchNullProbeResp	126045, 127045

Detecting a Block ACK DoS

The Block ACK mechanism that was introduced in 802.11e, and enhanced in 802.11nD3.0, has a built-in DoS vulnerability. The Bock ACK mechanism allows for a sender to use the ADDBA request frame to specify the sequence number window that the receiver should expect. The receiver will only accept frames in this window.

An attacker can spoof the ADDBA request frame causing the receiver to reset its sequence number window and thereby drop frames that do not fall in that range.

Detecting a ChopChop Attack

ChopChop is a plaintext recovery attack against WEP encrypted networks. It works by forcing the plaintext, one byte at a time, by truncating a captured frame and then trying all 256 possible values for the last byte with a corrected CRC. The correct guess causes the AP to retransmit the frame. When that happens, the frame is truncated again.

Detecting a Disconnect Station Attack

A disconnect attack can be launched in many ways; the end result is that the client is effectively and repeatedly disconnected from the AP.

Detecting an EAP Rate Anomaly

To authenticate wireless clients, WLANs may use 802.1x, which is based on a framework called Extensible Authentication Protocol (EAP). After an EAP packet exchange and the user is successfully authenticated, the EAP-Success is sent from the AP to the client. If the user fails to authenticate, an EAP-Failure is sent. In this attack, EAP-Failure or EAP-Success frames are spoofed from the access point to the client to disrupting the authentication state on the client. This confuses the client's state causing it to drop the AP connection. By continuously sending EAP Success or Failure messages, an attacker can effectively prevent the client from authenticating with the APs in the WLAN.

Detecting a FATA-Jack Attack Structure

FATA-Jack is an 802.11 client DoS tool that tries to disconnect targeted stations using spoofed authentication frames that contain an invalid authentication algorithm number.

Detecting a Hotspotter Attack

The Hotspotter attack is an evil-twin attack which attempts to lure a client to a malicious AP. Many enterprise employees use their laptop in Wi-Fi area hotspots at airports, cafes, malls etc. They have SSIDs of their hotspot service providers configured on their laptops. The SSIDs used by different hotspot service providers are well known. This enables the attackers to set up APs with hotspot SSIDs in close proximity of the enterprise premises. When the enterprise laptop Client probes for hotspot SSID, these malicious APs respond and invite the client to connect to them. When the client connects to a malicious AP, a number of security attacks can be launched on the client. A popular hacking tool used to launch these attacks is Airsnarf.

Detecting a Meiners Power Save DoS Attack

To save on power, wireless clients will "sleep" periodically, during which they cannot transmit or receive. A client indicates its intention to sleep by sending frames to the AP with the Power Management bit ON. The AP then begins buffering traffic bound for that client until it indicates that it is awake. An intruder could exploit this mechanism by sending (spoofed) frames to the AP on behalf of the client to trick the AP into believing the client is asleep. This will cause the AP to buffer most, if not all, frames destined for the client.

Detecting an Omerta Attack

Omerta is an 802.11 DoS tool that sends disassociation frames to all stations on a channel in response to data frames. The Omerta attack is characterized by disassociation frames with a reason code of 0x01. This reason code is “unspecified” and is not be used under normal circumstances.

Detecting Rate Anomalies

Many DoS attacks flood an AP or multiple APs with 802.11 management frames. These can include authenticate/associate frames which are designed to fill up the association table of an AP. Other management frame floods, such as probe request floods, can consume excess processing power on the AP.

Detecting a TKIP Replay Attack

TKIP is vulnerable to replay (via WMM/QoS) and plaintext discovery (via ChopChop). This affects all WPA-TKIP usage. By replaying a captured TKIP data frame on other QoS queues, an attacker can manipulate the RC4 data and checksum to derive the plaintext at a rate of one byte per minute.

By targeting an ARP frame and guessing the known payload, an attacker can extract the complete plaintext and MIC checksum. With the extracted MIC checksum, an attacker can reverse the MIC AP to Station key and sign future messages as MIC compliant, opening the door for more advanced attacks.

Detecting Unencrypted Valid Clients

An authorized (valid) client that is passing traffic in unencrypted mode is a security risk. An intruder can sniff unencrypted traffic (also known as *packet capture*) with software tools known as sniffers. These packets are then reassembled to produce the original message.

Detecting a Valid Client Misassociation

This feature does not detect attacks, but rather it monitors authorized (valid) wireless clients and their association within the network. Valid client misassociation is potentially dangerous to network security. The four types of misassociation that we monitor are:

- Authorized Client associated to Rogue—A valid client that is associated to a rogue AP
- Authorized Client associated to External AP—An external AP, in this context, is any AP that is not valid and not a rogue
- Authorized Client associated to Honeypot AP—A honeypot is an AP that is not *valid* but is using an SSID that has been designated as valid/protected
- Authorized Client in ad hoc connection mode—A valid client that has joined an ad hoc network

Detecting an AirJack Attack

AirJack is a suite of device drivers for 802.11(a/b/g) raw frame injection and reception. It was intended to be used as a development tool for all 802.11 applications that need to access the raw protocol, however one of the tools included allowed users to force off all users on an AP.

Detecting ASLEAP

ASLEAP is a tool created for Linux systems which is used to attack Cisco LEAP authentication protocol.

Detecting a Null Probe Response

A null probe response attack has the potential to crash or lock up the firmware of many 802.11 NICs. In this attack, a client probe-request frame will be answered by a probe response containing a null SSID. A number of popular NIC cards will lock up upon receiving such a probe response.

Configuring Intrusion Protection

Intrusion protection features support containment of an AP or a client. In the case of an AP, we will attempt to disconnect all client that are connected or attempting to connect to the AP. In the case of a client, the client's association to an AP is targeted. The following containment mechanisms are supported:

- Deauthentication containment: An AP or client is contained by disrupting its association on the wireless interface.
- Tarpit containment: An AP is contained by luring clients that are attempting to associate with it to a tarpit. The tarpit can be on the same channel as the AP being contained, or on a different channel (see Tarpit Shielding Overview on page 483).
- Wired containment: An AP or client is contained by disrupting its connection on the wired interface.

The WIP feature supports separate enforcement policies that use the underlying containment mechanisms to contain an AP or a client that do not conform to the policy. These policies are discussed in the sections that follow.

Understanding Infrastructure Intrusion Protection

Table 88 presents a summary of the infrastructure intrusion protection features with their related commands, traps, and syslog identifications. Details of each feature follow the table.

Table 88: *Infrastructure Protection Summary*

Feature	Command	Trap	Syslog ID
Protecting 40MHz 802.11 High Throughput Devices on page 473	`ids unauthorized-device-profile protect-ht-40mhz`	wlsxAPDeauthContainment wlsxClientDeauthContainment wlsxTarpitContainment	106005, 106006, 126102, 126103, 126108, 127102, 127103, 127108
Protecting 802.11n High Throughput Devices on page 473	`ids unauthorized-device-profile protect-high-throughput`	wlsxAPDeauthContainment wlsxClientDeauthContainment wlsxTarpitContainment	106005, 106006, 126102, 126103, 126108, 127102, 127103, 127108
Protecting Against Adhoc Networks on page 473	`ids unauthorized-device-profile protect-adhoc-network protect-adhoc-enhanced`	wlsxAPDeauthContainment wlsxClientDeauthContainment wlsxTarpitContainment wlsxEhancedAdhocContainment	106005, 106006, 126012, 126102, 126103, 126108, 127102, 127103, 127108, 126114
Protecting Against AP Impersonation on page 473	`ids impersonation-profile protect-ap-impersonation`	wlsxAPDeauthContainment wlsxClientDeauthContainment wlsxTarpitContainment	106005, 106006, 126102, 126103, 126108, 127102, 127103, 127108
Protecting Against Misconfigured APs on page 473	`ids unauthorized-device-profile protect-misconfigured-ap`	wlsxAPDeauthContainment wlsxClientDeauthContainment wlsxTarpitContainment	106005, 106006, 126102, 126103, 126108, 127102, 127103, 127108
Protecting SSIDs on page 474	`ids unauthorized-device-profile protect-ssid`	wlsxAPDeauthContainment wlsxClientDeauthContainment wlsxTarpitContainment	106005, 106006, 126102, 126103, 126108, 127102, 127103, 127108
Protecting Against Wireless Hosted Networks	ids unauthorized-device-profile detect-wireless-hosted-network protect-wireless-hosted-network	wlsxWirelessHostedNetwork-Detected wlsxClientAssociatedToHosted-NetworkDetected wlsxWirelessHostedNetwork-Containment wlsxHostOfWirelessNetwork-Containment	126110, 126111, 126112, 126113

Feature	Command	Trap	Syslog ID
Protecting Against Rogue Containment on page 474	`ids unauthorized-device-profile rogue-containment`	wlsxAPDeauthContainment wlsxClientDeauthContainment wlsxTarpitContainment	106005, 106006, 126102, 126103, 126108, 127102, 127103, 127108
Protecting Against Suspected Rogue Containment on page 474	`ids unauthorized-device-profile suspect-rogue-containment suspect-rogue-conf-level`	wlsxAPDeauthContainment wlsxClientDeauthContainment wlsxTarpitContainment	106005, 106006, 106010, 126102, 126103, 126108, 127102, 127103, 127108
Protection against Wired Rogue APs	ids general-profile wired-containment wired-containment-ap-adj-mac wired-containment-susp-l3-rogue	wlsxAPWiredContainment	126104,126105, 126106, 126107

Protecting 40MHz 802.11 High Throughput Devices

Protection from AP(s) that support 40MHz HT involves containing the AP such that clients can not connect.

Protecting 802.11n High Throughput Devices

Protection from AP(s) that support HT involves containing the AP such that clients can not connect.

Protecting Against Adhoc Networks

Protection from an adoc Network involves containing the adhoc network so that clients can not connect to it. The basic adhoc protection feature protects against adhoc networks using WPA/WPA2 security. The enhanced adhoc network protection feature protects against open/WEP adhoc networks. Both features can used together for maximum protection, or enabled or disabled separately



This feature requires that you enable the **wireless-containment** setting in the IDS general profile.

Protecting Against AP Impersonation

Protection from AP impersonation involves containing both the legitimate and impersonating AP so that clients can not connect to either AP.

Protecting Against Misconfigured APs

Protect Misconfigured AP enforces that valid APs are configured properly. An offending AP is contained by preventing clients from associating to it.

Protecting Against Wireless Hosted Networks

Clients using the Windows wireless hosted network feature can act as an access point to which other wireless clients can connect, effectively becoming a Wi-Fi HotSpot. This creates a security issue for enterprises, because unauthorized users can use a hosted network to gain access to the corporate network, and valid users that connect to a hosted network are vulnerable to attack or security breaches. This feature detects a wireless hosted network, and contains the client hosting this network.

Protecting SSIDs

Protect SSID enforces that valid/protected SSIDs are used only by valid APs. An offending AP is contained by preventing clients from associating to it.

Protecting Against Rogue Containment

By default, rogue APs are not automatically disabled. Rogue containment automatically disables a rogue AP by preventing clients from associating to it.

Protecting Against Suspected Rogue Containment

By default, suspected rogue APs are not automatically contained. In combination with the suspected rogue containment confidence level, suspected rogue containment automatically disables a suspect rogue by preventing clients from associating to it.

Protection against Wired Rogue APs

This feature enables containment from the wired side of the network.The basic wired containment feature in the IDS general profile isolates layer-3 APs whose wired interface MAC addresses are either the same as (or one character off from) their BSSIDs. The enhanced wired containment feature introduced in ArubaOS 6.3 can also identify and contain an AP with a preset wired MAC address that is completely different from the AP's BSSID. In many non-Dell APs, the MAC address the AP provides to wireless clients as a 'gateway MAC' is offset by one character from its wired MAC address.This enhanced feature allows ArubaOS to check to see if a suspected Layer-3 rogue AP's MAC address follows this common pattern.

Understanding Client Intrusion Protection

Table 89 list the client intrusion protection features with their related commands, traps, and syslog identifications. Details of each feature follow the table.

Table 89: *Client Protection Summary*

Feature	Command	Trap	Syslog ID
Protecting Valid Stations on page 474	`ids unauthorized-device-profile protect-valid-sta`	wlsxAPDeauthContainment wlsxClientDeauthContainment wlsxTarpitContainment	106005, 106006, 126102, 126103, 126108, 127102, 127103, 127108
Protecting Windows Bridge on page 475	`ids unauthorized-device-profile protect-windows-bridge`	wlsxAPDeauthContainment wlsxClientDeauthContainment wlsxTarpitContainment	106005, 106006, 126102, 126103, 126108, 127102, 127103, 127108

Protecting Valid Stations

Protecting a valid client involves disconnecting that client if it is associated to a non-valid AP.

Protecting Windows Bridge

Protecting from a Windows Bridge involves containing the client that is forming the bridge so that it can not connect to the AP.

Configuring the WLAN Management System (WMS)

The WLAN management system (WMS) on the controller monitors wireless traffic to detect any new AP or wireless client station that tries to connect to the network. When an AP or wireless client is detected, it is classified and its classification is used to determine the security policies which should be enforced on the AP or client.

In the WebUI

1. Navigate to the **Configuration > Advanced Services > Wireless** page.
2. Configure the parameters, as described in Table 90. Then click **Apply**.

Table 90: *WMS Configuration Parameters*

Parameter	Description
Adhoc AP Ageout	The amount of time, in minutes, that an adhoc (IBSS) AP unseen by any problems before it is deleted from the database. Enter 0 to disable ageout. Default: 30 minutes
AP Ageout Interval	The amount of time, in minutes, that an AP is unseen by any probes before it is deleted from the database. Enter 0 to disable ageout. Default: 30 minutes
AM Poll Interval	Interval, in milliseconds, for communication between the controller and Dell AMs. The controller contacts the AM at this interval to download AP to STA associations, update policy configuration changes, and download AP and STA statistics. Default: 60000 milliseconds (1 minute)
Number of AM Poll Retries	Maximum number of failed polling attempts before the polled AM is considered to be down. Default: 3
Station Ageout Interval	The amount of time, in minutes, that a client is unseen by any probes before it is deleted from the database. Enter 0 to disable ageout. Default: 30 minutes
Enable Statistics Update in DB	Enables or disables statistics update in the database. Default: enabled
Collect Stat	Enables collection of statistics (up to 25,000 entries) on the master controller for monitored APs and clients. This only applies when MMS is not configured. Default: disabled
Learn System Wired Mac	Enable or disable “learning” of wired MACs at the controller. Default: disabled
Propogate Wired Mac	Enables the propagation of the gateway wired MAC information.

Parameter	Description
	Default: enabled
Mark Neighbor APs as Persistent Neighbor APs	Enables or disables APs that are marked as neighbor from being aged out. Default: enabled
Learn APs	Enables or disables AP learning. Learning affects the way APs are classified. Default: disabled

In the CLI

Use the following commands to configure WMS via the CLI. The parameters in this command are described in detail in Table 90.

```
ids wms-general-profile
   adhoc-ap-ageout-interval <minutes> | ap-ageout-interval <minutes> | collect-stats
   {disable|enable} | learn-ap {enable|disable} | learn-system-wired-macs |
   persistent-neighbor {enable|disable} | poll-interval <milliseconds> |
   poll-retries <number> | propagate-wired-macs {enable|disable} | sta-ageout-interval
   <minutes> | stat-update {enable|disable}
```

Configuring Local WMS Settings

You can also use the CLI to define local WMS system settings for the maximum number of APs and client stations.



Use this command with caution. Increasing the limit will cause an increase in usage in the memory by WMS. In general, each entry will consume about 500 bytes of memory. If the setting is bumped up by 2000, then it will cause an increase in WMS memory usage by 1MB

```
(host) (config) #ids wms-local-system-profile max-threshold <max-threshold>
```

Managing the WMS Database

The WMS process interacts with all the air monitor (AM) processes in the network. When WMS receives an event message from an AM, the WMS process will save the event information along with the BSSID of the AP that generated the event in the WMS database. Use the following commands in **Enable** mode to manage the WMS database.

The **wms export-db** command exports the specified file as an ASCII text file into the WMS database.

```
(host) #wms export-db database <file>
```

The **wms import-db** command imports the specified file into the WMS database:

```
(host) #wms import-db database <file>
```

The **wms reint-db** command reinitializes the WMS database. Note that this command does not make an automatic backup of the current database.

```
(host) #wms reint-db
```

Understanding Client Blacklisting

When a client is blacklisted in the Dell system, the client is not allowed to associate with any AP in the network for a specified amount of time. If a client is connected to the network when it is blacklisted, a deauthentication message is sent to force the client to disconnect. While blacklisted, the client cannot associate with another SSID in the network.

The controller retains the client blacklist in the user database, so the information is not lost if the controller reboots. When you import or export the controller's user database, the client blacklist will be exported or imported as well.

Methods of Blacklisting

There are several ways in which a client can be blacklisted in the Dell system:

- You can manually blacklist a specific client. See Blacklisting Manually on page 477 for more information.
- A client fails to successfully authenticate for a configured number of times for a specified authentication method. The client is automatically blacklisted. See Blacklisting by Authentication Failure on page 477 for more information.
- A DoS or man in the middle (MITM) attack has been launched in the network. Detection of these attacks can cause the immediate blacklisting of a client. See Enabling Attack Blacklisting on page 478 for more information.
- An external application or appliance that provides network services, such as virus protection or intrusion detection, can blacklist a client and send the blacklisting information to the controller via an XML API server. When the controller receives the client blacklist request from the server, it blacklists the client, logs an event, and sends an SNMP trap.

  See External Services Interface on page 992 for more information.



The External Services Interface feature require the Policy Enforcement Firewall Next Generation (PEFNG) license installed in the controller.

Blacklisting Manually

There are several reasons why you may choose to blacklist a client. For example, you can enable different Dell intrusion detection system (IDS) features that detect suspicious activities, such as MAC address spoofing or DoS attacks. When these activities are detected, an event is logged and an SNMP trap is sent with the client information. To blacklist a client, you need to know its MAC address.

To manually blacklist a client via the WebUI:

1. Navigate to the **Monitoring > Controller > Clients** page.
2. Select the client to be blacklisted and click the **Blacklist** button.

To clear the entire client blacklist using the WebUI:

1. Navigate to the **Monitoring > Controller > Clients** page.
2. Click **Remove All from Blacklist**.

To manually blacklist a client via the command-line interface, access the CLI in config mode and issue the following command:

```
stm add-blacklist-client <macaddr>
```

To clear the entire client blacklist using the command-line interface, access the CLI in config mode and issue the following command:

```
stm purge-blacklist-client
```

Blacklisting by Authentication Failure

You can configure a maximum authentication failure threshold for each of the following authentication methods:

- 802.1x
- MAC

- Captive portal
- VPN

When a client exceeds the configured threshold for one of the above methods, the client is automatically blacklisted by the controller, an event is logged, and an SNMP trap is sent. By default, the maximum authentication failure threshold is set to 0 for the above authentication methods, which means that there is no limit to the number of times a client can attempt to authenticate.

With 802.1x authentication, you can also configure blacklisting of clients who fail machine authentication.



When clients are blacklisted because they exceed the authentication failure threshold, they are blacklisted indefinitely by default. You can configure the duration of the blacklisting; see Setting Blacklist Duration on page 479.

To set the authentication failure threshold via the WebUI:

1. Navigate to the **Configuration > Security > Authentication > Profiles** page.
2. In the **Profiles** list, select the appropriate authentication profile, then select the profile instance.
3. Enter a value in the **Max Authentication failures** field.
4. Click **Apply**.

To set the authentication failure threshold via the command-line interface, access the CLI in config mode and issue the following commands:

```
aaa authentication {captive-portal|dot1x|mac|vpn} <profile>
   max-authentication-failures <number>
```

Enabling Attack Blacklisting

There are two type of automatic client blacklisting that can be enabled: blacklisting due to spoofed deauthentication, or blacklisting due to other types of DoS attacks.

Automatic blacklisting for DoS attacks other than spoofed deauthentication is enabled by default. You can disable this blacklisting on a per-SSID basis in the virtual AP profile.

Man in the middle (MITM) attacks begin with an intruder impersonating a valid enterprise AP. If an AP needs to reboot, it sends deauthentication packets to connected clients to enable them to disconnect and reassociate with another AP. An intruder or attacker can spoof deauthentication packets, forcing clients to disconnect from the network and reassociate with the attacker's AP. A valid enterprise client associates to the intruder's AP, while the intruder then associates to the enterprise AP. Communication between the network and the client flows through the intruder (the man in the middle), thus allowing the intruder the ability to add, delete, or modify data. When this type of attack is identified by the Dell system, the client can be blacklisted, blocking the MITM attack. Enable this blacklisting ability in the IDS DoS profile (this is disabled by default).

To enable spoofed deauth detection and blacklisting via the WebUI:

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. Select either **AP Group** or **AP Specific** tab. Click **Edit** for the AP group or AP name.
3. In the Profiles list, expand the **IDS** menu, then select **IDS profile**.
4. Select the **IDS DOS profile**.
5. Select (check) **Spoofed Deauth Blacklist**.
6. Click **Apply**.

To enabled spoofed deauth detection and blacklisting via the command-line interface, access the CLI in config mode, and issue the following commands:

```
ids dos-profile <profile>
   spoofed-deauth-blacklist
```

Setting Blacklist Duration

You can configure the duration that clients are blacklisted on a per-SSID basis via the virtual AP profile. There are two different blacklist duration settings:

- For clients that are blacklisted due to authentication failure. By default, this is set to 0 (the client is blacklisted indefinitely).
- For clients that are blacklisted due to other reasons, including manual blacklisting. By default, this is set to 3600 seconds (one hour). You can set this to 0 to blacklist clients indefinitely.

To configure the blacklist duration via the WebUI:

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. Select either AP Group or AP Specific tab. Click **Edit** for the AP group or AP name.
3. In the Profiles list, select **Wireless LAN**, then **Virtual AP**. Select the virtual AP instance.
   - To set a blacklist duration for authentication failure, enter a value for **Authentication Failure Blacklist Time**.
   - To set a blacklist duration for other reasons, enter a value for **Blacklist Time**.
4. Click **Apply**.

To configure the blacklist duration via the command-line interface, access the CLI in config mode and issue the following commands:

```
wlan virtual-ap <profile>
   auth-failure-blacklist-time <seconds>
   blacklist-time <seconds>
```

Removing a Client from Blacklisting

You can manually remove a client from blacklisting using either the WebUI or CLI:

To remove a client from blacklisting via the WebUI:

1. Navigate to the **Monitoring > Controller > Blacklist Clients** page.
2. Select the client that you want to remove from the blacklist, then click **Remove from Blacklist**.

To remove a client from blacklisting via the command-line interface, access the CLI in enable mode and issue the following command:

```
stm remove-blacklist-client <macaddr>
```

Working with WIP Advanced Features

Device Classification is the first step in securing the corporate environment from unauthorized wireless access. Adequate measures that quickly shut down intrusions are critical in protecting sensitive information and network resources. APs and stations must be accurately classified to determine whether they are valid, rogue, or a neighboring AP. Then, an automated response can be implemented to prevent possible intrusion attempts.

TotalWatch allows for detecting devices that are running on typical operational channels. Tarpit Shielding provides a better way of containing devices that are deemed unauthorized. Both of these features are discussed in the sections that follow.

- Configuring TotalWatch on page 480
- Administering TotalWatch on page 482
- Tarpit Shielding Overview on page 483
- Configuring Tarpit Shielding on page 484

Configuring TotalWatch

Dell 802.11n APs and non-11n APs in AM-mode support for TotalWatch is the ability to scan all channels of the RF spectrum, including 2.4-and 5-GHz bands as well as the 4.9-GHz public safety band. TotalWatch also provides 5-MHz granular channel scanning of bands for rogue devices, and dynamic scanning dwell times to focus on those channels with traffic. TotalWatch provides an advanced set of features to detect unauthorized wireless devices and a set of customized rules are used to highlight devices that truly pose a threat to the network.



TotalWatch is supported on APs deployed in the AM-mode only.

TotalWatch provides monitoring support for the entire WLAN spectrum. Dell APs in the AM-mode can *monitor* the following frequencies:

- 2412MHz to 2472MHz in the 2.5GHz band
- 5100Mhz to 5895MHz in the 5GHz band.

Dell APs in AM-mode can *scan* the following additional frequencies:

- 2484 MHz and 4900Mhz to 5000MHz (J-channels)
- 5000 to 5100Mhz

If the AP is HT-capable (High Throughput), then these frequencies are scanned in the 40MHz mode.

Understanding TotalWatch Channel Types and Qualifiers

Based on the regulatory characteristics, channels are categorized into the following types:

Reg-domain Channels—A channel that belongs to the regulatory domain of the country in which the AP is deployed. The set of channels that belong to this group is a subset of the channels in all-reg-domain channel group.

All-reg-domain Channels—A valid non-overlapping channel that is in the regulatory domain of at least one country. The channels in this category belong to the frequency range of:

- 2412MHz to 2472MHz in the g-band
- 5100Mhz to 5895MHz in the a-band.

Rare Channel—Channels that fall into a frequency range outside of the regulatory domain; 2484 MHz and 4900MHz-4995MHz (J-channels), and 5000-5100Mhz. The channels in this group do not belong to any other group.

Each of these channel types can have an associated qualifier:

Active Channel—This qualifier indicates that wireless activity is detected on this channel by the presence of an AP or other 802.11 activity; a probe requests for example.

DOS Channel—A channel where wireless containment is active. This channel should belong to the country-code channel (regulatory domain).

Understanding TotalWatch Monitoring Features

TotalWatch enables monitoring of all channels including regulatory domain and rare channels. You can select one of the following scanning modes for each radio AP.

- scan only the channels that belong to the AP's regulatory domain
- scan channels that belong to all regulatory domains
- scan all channels

Understanding TotalWatch Scanning Spectrum Features

TotalWatch scans the following frequencies.

- G-band—2412MHz to 2472MHz
- J band—2484 MHz
- A-band—5000-5100Mhz to 5895MHz
- J-band—4900-4995MHz

Table 91 list the frequency-to-channel mapping used by TotalWatch.

Table 91: *Frequency to Channel Mapping*

Frequency	Channel
2412 – 2472MHz (in increments of 5MHz)	1 - 13
2484MHz	14
5100 – 5895MHz (in increments of 5MHz)	20 - 179
4900 – 4995MHz (in increments of 5MHz)	180 - 199
5000 – 5100MHz	200 - 219

Understanding TotalWatch Channel Dwell Time

When an AP (in am-mode) visits a channel, the amount of time the AP *stays* on that channel is known as the *dwell time*. The channel dwell time is a variable value based on the following channel types.

dwell-time-active-channel—For channels where there is wireless activity. Default setting is 500 ms.

dwell-time-reg-domain channel—For channels that belong to the AP's regulatory domain group (reg-domain) with *no* wireless activity. The default setting is 250 ms.

dwell-time-other-reg-domain-channel—For channels that belong to the *all* regulatory domain group (all-reg-domain) with *no* wireless activity The default setting is 250 ms.

dwell-time-rare-channel—For channels in the rare group where *no* wireless activity is detected. The default value is 100 ms.

Use the **rf am-scan-profile** command to set the dwell time and scan mode.

Understanding TotalWatch Channel Visiting

The Active and DOS channels are visited more frequently than the other channels. The order of preference in selecting the next channel is:

1. DOS
2. Active
3. reg-domain
4. All-reg-domain
5. Rare

Once a channel is selected, the dwell time for that channel is determined based on the channel type. At the end of the dwell time, a new channel is picked.

Understanding TotalWatch Age out of Devices

ArubaOS uses a combination of inactivity time and unseen time to age out a device. This ensures that the channel is scanned a sufficient number of times before a device ages out. AM module maintains the following parameters:

Discovered Time—The absolute time, in seconds, since the device was discovered.

Monitored Time—The number of times the channel was scanned since discovery.

Inactivity Time—The number of times the device was not "seen" when the channel is scanned.

Unseen Time—The absolute time, in seconds, since the device was last "seen."

Administering TotalWatch

The AM module will initialize the channel list for each of the AP's radio based on the scan mode setting for the radio. For example, if scan mode is set to rare, then the channel list will contain all possible channels. You can view these channels by using the **show ap arm scan-times** command.

Configuring Per Radio Settings

For each radio, you can configure the following settings (for detailed information on commands, refer to the *Dell Networking W-Series ArubaOS 6.4.x Command Line Reference Guide*):

- the dwell times for the various channel types
- the channel list that should be used for scanning

These settings are configured via the command **rfam-scan-profile**, which can be attached to the two profiles, **dot11a-radio-profile** and **dot11g-radio-profile**.

The am-scan-profile includes the following parameters that can be configured:

```
rf am-scan-profile <name>
scan-mode [reg-domain | all-reg-domain | rare]
```

The default setting is the all-reg-domain. This is consistent with the default functioning of the AM scanning where the radio scans channels belonging to all regulatory domains.

Configuring Per AP Setting

If the AP is a dual-band single radio AP, an option is available to specify which band should be used for scanning in AM-mode. This setting is available in the "ap system-profile", via the am-scan-rf-band command.

```
ap system-profile <name>
am-scan-rf-band [a | g | all]
```

The default value is "all", which is consistent with the prior behavior. This setting is ignored in the case of a dual radio AP.

There are four parameters that will control the age out of devices in the AM module.

```
ids general-profile <name>
ap-inactivity-timeout
sta-inactivity-timeout
ap-max-unseen-timeout
sta-max-unseen-timeout
```

The inactivity timeout is the number of times the device was not "seen" when the channel was scanned. The unseen timeout is the time, in seconds, since the device was last "seen."

The **show ap monitor scan-info/channel** commands provide details of the channel types, dwell times, and the channel visit sequence.

```
(host) # show ap monitor scan-info ap-name rb-121
```

```
WIF Scanning State: wifi0
-------------------------
Parameter                      Value
---------                      -----
Scan Mode                      all-reg-domain
Scan Channel                   yes
Disable Scanning               no
Current Channel                36-
Current Scan Channel           36-
Current Channel Index          1
Current Scan Start Milli Tick  351757100
Current Dwell Time             600
Current Scan Type              active

Scan-Type-Info
--------------
Info-Type          Active  Reg-domain  All-reg-domain  Rare  DOS
---------          ------  ----------  --------------  ----  ---
Dwell Times        600     250         100             100   600
Last Scan Channel  36-     116-        0               0     0

(host) #show ap monitor channel ap-name rb-121 36

Aggregate Stats
---------------
retry  low-speed  non-unicast  frag  bwidth  phy-err  mac-err  noise
-----  ---------  -----------  ----  ------  -------  -------  -----
0      0          0            0     1       0        9        90

Scanning Stats
--------------
scans-attempted  assign-time(ms)  last-visit-time  monitored-time  reside-time(ms)  dos-scans
flags
---------------  ---------------  ---------------  --------------  ---------------  ---------
-----
42702              25620500         402424           56245           0
    0          DVACLU
Channel Flags: D: Default, V: Valid, A: AP Present, C: Reg Domain Channel,
               O: DOS Channel, Z: Rare Channel
               T: Valid 20MHZ Channel, F: Valid 40MHz Channel,
               L: Scan 40MHz Channel (lower), U: Scan 40MHz channel (upper)
               R: Radar detected in last 30 min, X: DFS required
```

Licensing

The ability to perform rare scanning is available only with the RFprotect license. However, the AP can scan 'reg-domain' or 'all-reg-domain' channels without the RFprotect license.

Tarpit Shielding Overview

The Tarpit Shielding feature is a type of wireless containment. Detected devices that are classified as rogues are contained by forcing client association to a fake channel or BSSID. This method of tarpitting is more efficient than rogue containment via repeated de-authorization requests. Tarpit Sheilding works by spoofing frames from an AP to *confuse* a client about its association. The *confused* client assumes it is associated to the AP on a different (fake) channel than the channel that the AP is actually operating on, and will attempt to communicate with the AP in the fake channel.

Tarpit Shielding works in conjunction with the *deauth* wireless containment mechanism. The deauth mechanism triggers the client to generate probe request and subsequent association request frames. The AP then responds with probe response and association response frames. Once the monitoring AP sees these frames, it will spoof the probe-response and association response frames, and manipulates the content of the frames to confuse the client.

A station is determined to be in the Tarpit when we *see* it sending data frames in the fake channel. With some clients, the station remains in tarpit state until the user manually disables and re-enables the wireless interface.

Configuring Tarpit Shielding

Tarpit shielding is configured on an AP using one of two methods:

Disable all clients—In this method, any client that attempts to associate with an AP marked for containment is sent spoofed frames.

Disable non-valid clients—In this method, only non-authorized clients that attempt to associate with an AP is sent to the tarpit.

The choices for disabling Tarpit Shielding on an AP are:

- Deauth-wireless-containment
- Deauth-wireless-containment with tarpit-shielding (excluding-valid-clients)
- Deauth-wireless-containment with tarpit-shielding

EnablingTarpit Shielding

Use the **ids-general-profile** command to configure Tarpit Shielding (for detailed information on commands refer to the *Command Line Reference Guide*).

```
ids general-profile default
   wireless-containment [deauth-only | none | tarpit-all-sta | tarpit-non-valid-sta]
```

Use the following show commands to view updated Tarpit Shielding status and the spoofed frames generated for an AP:

```
show ap monitor stats …
show ap monitor containment-info
```

Understanding Tarpit Shielding Licensing CLI Commands

In the **ids general-profile default wireless-containment** command, the 'tarpit-non-valid-sta' and 'tarpit-all-sta' options are available only with a RFprotect license. The 'deauth-only' and 'none' options are available with the Base OS license.

In ArubaOS, related configuration parameters are grouped into *profiles* that you can apply as needed to an AP group or to individual APs. When an AP is first installed on the network and powered on, the AP locates its host controller and the AP's designated configuration is "pushed" from the controller to the AP. This chapter gives an overview of the basic function of each AP profile, and describes the process to install and configure the APs on your network.

The following topics are included in this chapter:

- Basic Functions and Features on page 485
- Naming and Grouping APs
- Understanding AP Configuration Profiles on page 488
- Before you Deploy an AP on page 495
- Enable Controller Discovery on page 495
- Enable DHCP to Provide APs with IP Addresses
- Enable Controller Discovery on page 495
- Configuring Installed APs on page 500
- Optional AP Configuration Settings on page 503
- Configuring AP Image Preload on page 790
- RF Management on page 509
- Optimizing APs Over Low-Speed Links on page 521
- Configuring AP Channel Assignments on page 526
- Managing AP Console Settings on page 528
- Link Aggregation Support on W-AP220 Series and W-AP270 Series on page 529
- Service Tag on page 531

Basic Functions and Features

You configure APs using the WebUI and the CLI on the controller. Table 92 list the basic configuration functions and features.

Table 92: *AP Configuration Function Overview*

Features and Function	Description
Wireless LANs	A wireless LAN (WLAN) permits wireless clients to connect to the network. An AP broadcasts the SSID (which corresponds to a WLAN configured on the controller) to wireless clients. APs support multiple SSIDs. WLAN configuration includes the authentication method and the authentication servers by which wireless users are validated for access. The WebUI includes a WLAN Wizard that provides easy-to-follow steps to configure a new WLAN. **NOTE:** All new WLANs are associated with the ap-group named “default”.
AP operation	An AP can function as an AP that serves clients, as an air monitor (AM) performing network and radio frequency (RF) monitoring, or as a hybrid AP that serves both clients and performs spectrum analysis a single radio channel. You can also specify the regulatory domain (the country) which determines the 802.11 transmission spectrum in which the AP will operate. Within the regulated transmission spectrum, you can configure 802.11a, 802.11b/g, or 802.11n (high-throughput) radio settings. **NOTE:** The 802.11n features, such as high-throughput and 40 MHz configuration settings, are supported on APs that are 802.11n standard compliant.
Quality of Service (QoS)	Configure Voice over IP call admission control options and bandwidth allocation for 5 GHz (802.11a) or 2.4 GHz (802.11b/g) frequency bands of traffic.
RF Management	Configure settings for balancing wireless traffic across APs, detect holes in radio coverage, or other metrics that can indicate interference and potential problems on the wireless network. Adaptive Radio Management (ARM) is an RF spectrum management technology that allows each AP to determine the best 802.11 channel and transmit power settings. ARM provides several configurable settings.
Intrusion Detection System	Configure settings to detect and disable rogue APs, ad-hoc networks, and unauthorized devices, and prevent attacks on the network. You can also configure signatures to detect and prevent intrusions and attacks.
Mesh	Configure Dell APs as mesh nodes to bridge multiple Ethernet LANs or extend wireless coverage. A mesh node is either • a mesh portal: an AP that uses its wired interface to reach the controller • a mesh point:an AP that establishes a path to the controller via the mesh portal Mesh environments use a wireless backhaul to carry traffic between mesh nodes. This allows one 802.11 radio to carry traditional WLAN services to clients and one 802.11 radio to carry mesh traffic and WLAN services. Secure Enterprise Mesh on page 533 contains more specific information on the Mesh feature.

Naming and Grouping APs

In the Dell user-centric network, each AP has a unique name and belongs to an AP group.

Each AP is identified with an automatically-derived name. The default name depends on if the AP has been previously configured.

- The AP has not been configured—the name is the AP’s Ethernet MAC address in colon-separated hexadecimal digits.
- Configured with a previous ArubaOS release—the name is in the format *building.floor.location*

You can assign a new name (up to 63 characters) to an AP; the new name must be unique within your network. For example, you can rename an AP to reflect its physical location within your network, such as “building3-lobby”.

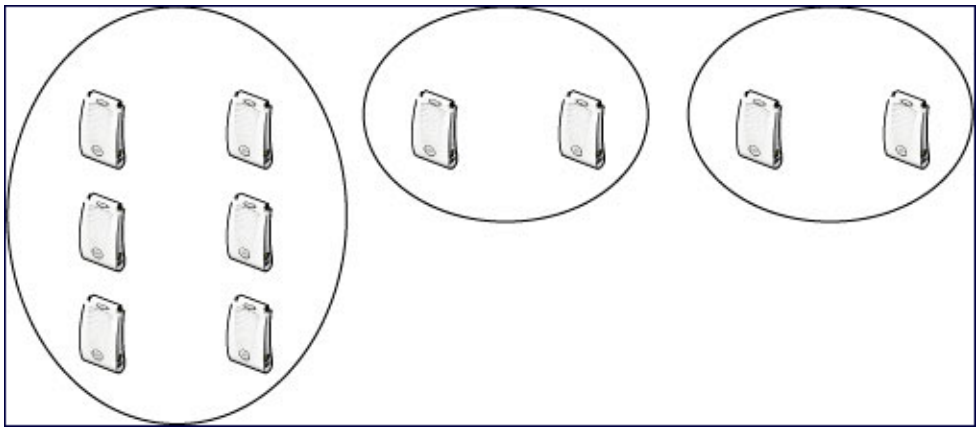


Renaming an AP requires a reboot of the AP before the new name takes effect. Therefore, wait until there is little or no client traffic passing through the AP before renaming it.

An *AP group* is a set of APs to which the same configuration is applied. There is an AP group called “default” to which all APs discovered by the controller are assigned. By using the “default” AP group, you can configure features that are applied globally to all APs.

You can create additional AP groups and assign APs to that new group. However, an AP can belong to only one AP group at a time. For example, you can create an AP group “Victoria” that consists of the APs that are installed in a company’s location in British Columbia. You can create another AP group “Toronto” that consists of the APs in Ontario. You can configure the “Toronto” AP group with different information from the APs in the “Victoria” AP group (see Figure 50).

Figure 50 *AP Groups*



While you can use an AP group to apply a feature to a set of APs, you can also configure a feature or option for a specific AP by referencing the AP’s name. Any options or values that you configure for a specific AP will override the same options or values configured for the AP group to which the AP belongs.

The following procedures describes how to create an AP group.



Reassigning an AP from an AP group requires a reboot of the AP for the new group assignment to take effect. Therefore,wait until there is little or no client traffic passing through the AP before reassigning it.

Creating an AP group

You can use the WebUI or the CLI to create a new AP group.

In the WebUI

1. Navigate to the **Configuration > Wireless> AP Configuration > AP Group** page.
2. Click **New**. Enter the new AP group name and click **Add**. The new AP group appears in the Profile list.

In the CLI

Use the following command to create an AP group:

```
ap-group <group>
```

When you create an AP group with the CLI, you can specify the virtual AP definitions and configuration profiles you want applied to the APs in the group.

Assigning APs to an AP Group

Although you will assign an AP to an AP group when you first deploy the device, you can assign an AP to a different AP group at any time.



Once the **ap-regroup** command is executed, the AP automatically reboots. If the AP is powered off or otherwise not connected to the network or controller, the executed command is queued until the AP is powered on or reconnected. Again, the AP will automatically reboot as soon as the command is executed.

In the WebUI

1. Navigate to the **Configuration > Wireless> AP Installation** page. The list of discovered APs appears in this page (all discovered APs initially belong to the AP group named "default").
2. Select the AP you want to reassign, and click **Provision**. From the Provisioning page, select the AP group from the drop-down menu.
3. Click **Apply and Reboot**.

In the CLI

Use the following command to assign a single AP to an existing AP group. Use the WebUI to assign multiple APs to an AP group at the same time.

```
ap-regroup {ap-name <name>|serial-num <number>|wired-mac <macaddr>} <group>
```

Understanding AP Configuration Profiles

An AP configuration profile is a general name to describe any of the different groups of settings that can defined, saved, and applied to an Access Point. ArubaOS has many different types of profiles that each allow you to configure a different aspect of an AP's overall configuration. ArubaOS also contains a predefined "default" profile for each profile type. You can use the predefined settings in these default profiles, or create entirely new profiles that you can edit as required.

Each different AP configuration profile type can be managed using the CLI or the WebUI. To see a full list of available configuration profiles using the command-line interface, access the CLI and issue the command **show profile-hierarchy**. To view available configuration profiles using the WebUI, select the **Configuration** tab in the and navigate to **Advanced Services > All Profiles**.

The **All Profiles** tab arranges the different AP configuration profile types into the following categories:

- AP Profiles
- RF Management Profiles
- Wireless LAN Profiles
- Mesh Profiles
- QoS Profiles
- IDS Profiles
- HA Group profiles
- Other Profiles



The profile types that appear in the **All Profiles** tab may vary, depending upon the controller configuration and available licenses.

AP Profiles

The AP profiles configure AP operation parameters, radio settings, port operations, regulatory domain, and SNMP information.

- **AP system profile**: defines administrative options for the controller, including the IP addresses of the local, backup, and master controllers, Real-time Locating Systems (RTLS) server values and the number of consecutive missed heartbeats on a GRE tunnel before an AP reboots. For details on configuring this profile, see Optional AP Configuration Settings.
- **Regulatory domain**: defines the AP's country code and valid channels for both legacy and high-throughput 802.11a and 802.11b/g radios. For examples on figuring a regulatory domain profile, see Configuring AP Channel Assignments on page 526.
- **Wired AP profile**: determines if 802.11 frames are tunneled to the controller using Generic Routing Encapsulation (GRE) tunnels, bridged into the local Ethernet LAN, or configured for a combination of the two (split-mode). In tunnel forwarding mode, the AP handles all 802.11 association requests and responses, but sends all 802.11 data packets, action frames and EAPOL frames over a GRE tunnel to the controller for processing. When a remote AP or campus AP is in bridge mode, the AP handles all 802.11 association requests and responses, encryption/decryption processes, and firewall enforcement. In split-tunnel mode, 802.11 frames are either tunneled or bridged, depending on the destination (corporate traffic goes to the controller, and Internet access remains local). For details, see Configuring Ethernet Ports for Mesh on page 561
- **AP LLDP-MED Network Policy and AP LLDP profiles**:link Layer Discovery Protocol (LLDP), is a Layer-2 protocol that allows network devices to advertise their identity and capabilities on a LAN. The LLDP-MED Network Policy profile defines the VLAN, priority levels, and DSCP values used by a voice or video application. Wired interfaces on Dell APs support LLDP by periodically transmitting LLDP Protocol Data Units (PDUs) comprised of selected type-length-value (TLV) elements. The AP LLDP profile identifies which TLVs will be sent by the AP. For details, see Understanding Extended Voice and Video Features on page 915.
- **Ethernet interface profile**:sets the duplex mode and speed of the AP's Ethernet link. The configurable speed is dependent on the port type, and you can define a separate Ethernet Interface profile for each Ethernet link. For details on configuring this profile, see Table 95.
- **Ethernet Interface Port/Wired Port Profile:** specifies a AAA profile for users connected to the wired port on an AP. For details on configuring this profile, see Securing Clients on an AP Wired Port on page 868
- **AP Provisioning profile**: defines a group of provisioning parameters for an AP or AP group. For details on configuring this profile, see .
- **AP Authorization Profile**—Allows you to assign an to a provisioned but unauthorized AP to a AP group with a restricted configuration profile. For details see Configuring Remote AP Authorization Profiles on page 655.
- **EDCA parameters profile (Station)**:client to AP traffic prioritization parameters, including Enhanced Distributed Channel Access (EDCA) parameters for background, best-effort, voice and video queues. For additional information on configuring this profile, see Using the WebUI to configure EDCA parameters on page 897.
- **EDCA parameters profile (AP):** AP to client traffic prioritization, including EDCA parameters for background, best-effort, voice and video queues. For additional information on configuring this profile, see Using the WebUI to configure EDCA parameters on page 897.
- **Spectrum Local Override Profile**:configure an individual AP radio as a spectrum monitor, For details, see Converting an Individual AP to a Spectrum Monitor on page 706.

RF Management Profiles

The profiles configure radio tuning and calibration, AP load balancing, and RSSI metrics.

- **802.11a radio profile:** defines AP radio settings for the 5 GHz frequency band, including the Adaptive Radio Management (ARM) profile and the high-throughput (802.11n) radio profile. For additional information on configuring this profile, see 802.11a and 802.11g RF Management Profiles on page 509.
- **802.11g radio profile:** defines AP radio settings for the 2.4 GHz frequency band, including the Adaptive Radio Management (ARM) profile and the high-throughput (802.11n) radio profile. Each 802.11a and 802.11b radio profile includes a reference to an Adaptive Radio Management (ARM) profile.

  If you want the ARM feature to dynamically select the best channel and transmission power for the radio, verify that the 802.11a/802.11g radio profile references an active and enabled ARM profile. If you want to manually select a channel for each AP group, create separate 802.11a and 802.11g profiles for each AP group and assign a different transmission channel for each profile. For additional information on configuring this profile, see 802.11a and 802.11g RF Management Profiles on page 509.
- **ARM profile:** defines the Adaptive Radio Management (ARM) settings for scanning, acceptable coverage levels, transmission power and noise thresholds. In most network environments, ARM does not need any adjustments from its factory-configured settings. However, if you are using VoIP or have unusually high security requirements you may want to manually adjust the ARM thresholds. For complete details on Adaptive Radio Management, refer to Adaptive Radio Management (ARM) on page 432.
- **High-throughput radio profile:** manages high-throughput (802.11n) radio settings for 802.11n-capable APs. A high-throughput profile determines 40 Mhz tolerance settings, and controls whether or not the APs using this profile will advertise intolerance of 40 MHz operation. (This option is disabled by default, allowing 40 MHz operation.) For additional information on configuring this profile, see High-Throughput Virtual APs on page 424.
- **RF Optimization profile:** enables or disables load balancing based on a user-defined number of clients or degree of AP utilization on an AP. Use this profile to detect coverage holes, radio interference and STA association failures and configure Received signal strength indication (RSSI) metrics.
- **RF Event Thresholds profile:** defines error event conditions, based on a customizable percentage of low-speed frames, non-unicast frames, or fragmented, retry or error frames. For additional information on configuring this profile, see RF Event Configuration on page 519.
- **AM Scanning**: Dell 802.11n APs and non-11n APs in AM-mode support the TotalWatch scanning feature giving them the ability to scan all channels of the RF spectrum, including 2.4-and 5-GHz bands as well as the 4.9-GHz public safety band. The AM Scanning profile enables this feature, and defines the dwell types for different channel types.

Wireless LAN Profiles

The Wireless LAN collection of profiles configure WLANs in the form of virtual AP profiles. A virtual AP profile contains an SSID profile which defines the WLAN, the high-throughput SSID profile, and an AAA profile that defines the authentication for the WLAN.

Unlike other profile types, you can configure and apply multiple instances of virtual AP profiles to an AP group or to an individual AP.

- **802.11k profile:** manages settings for the 802.11k protocol. The 802.11k protocol allows APs and clients to dynamically query their radio environment and take appropriate connection actions. For example: In a 802.11k network if the AP with the strongest signal reaches its CAC (Call Admission Control) limits for voice calls, then *on-hook* voice clients may connect to an under utilized AP with a weaker signal. You can configure the following options in 802.11k profile:
  - Enable or disable 802.11K support on the AP
  - Forceful disassociation of on-hook voice clients
  - Measurement mode for beacon reports.

  For more details, see Radio Resource Management (802.11k) on page 404.

- **Handover Trigger profile**: configure a handover trigger profile to ensure QoS for voice calls for APs with the 802.11k feature enabled. For more details, see Enabling Wi-Fi Edge Detection and Handover for Voice Clients on page 922
- **RRM IE profile**:configure a Radio Resource Management Information Element (RRM IE) profile to define the information elements advertised by an AP with 802.11k support enabled. For more details, see Configuring Radio Resource Management Information Elements on page 406
- **Beacon Report Request Profile:** APs with the 802.11k feature enabled use request messages to solicit measurements. This profile defines the information an AP can send in beacon report requests. For details, see Understanding AP Configuration Profiles on page 488
- **802.11r profile:** APs with the 802.11r (Fast BSS Transition) feature enabled minimize the delay when a client transitions from one BSS to another within the same ESS. For more details, see Fast BSS Transition (802.11r) on page 412
- **TSM Report Request Profile**: APs with the 802.11k feature enabled use request messages to solicit measurements. This profile defines the information an AP can send in traffic stream measurement reports. For more details, see Understanding AP Configuration Profiles on page 488
- **SSID profile:** Configures network authentication and encryption types. This profile also includes references to the EDCA (enhanced distributed channel access) Parameters Station Profile, the EDCA Parameters AP Profile and a High-throughput SSID profile.

  Use this profile to configure basic settings such as 802.11 authentication and encryption settings, or advanced settings such as DTIM (delivery traffic indication message) intervals, 802.11a/802.11g basic and transmit rates, DHCP settings and WEP keys. The advanced SSID profile settings allows you to deny broadcast probes and hide the SSID. For details on configuring an SSID profile, see SSID Profiles on page 414.



Beacon rates for 802.11a and 802.11g beacons should only be configured on APs with Distributed Antenna Systems (DAS). Configuring beacon rates during normal operation may cause connectivity problems.

- **High-throughput SSID profile:** high-throughput APs support additional settings not available in legacy APs. A High-throughput SSID profile enables/disables high-throughput (802.11n) features with 40 MHz channel usage, and define values for aggregated MAC protocol data units (MDPUs) and Modulation and Coding Scheme (MCS) ranges. If you modify a currently provisioned and running high-throughput SSID profile, your changes take effect immediately; rebooting is not required. For details on configuring a high-throughput SSID profile, see High-Throughput Virtual APs on page 424.
- **Virtual AP profile:** this profile defines your WLAN by enabling or disabling the band steering, fast roaming and DoS prevention features. It defines radio band, forwarding mode and blacklisting parameters, and includes references to an AAA Profile, 802.11K Profile, and a High-throughput SSID profile. You can apply multiple virtual AP profiles to an AP group or to an individual AP; for most other profiles, you can apply only one instance of the profile to an AP group or AP at a time. For details on configuring a Virtual AP profile, see Virtual AP Profiles on page 395.
- **VIA Client WLAN profile**:the VIA client WLAN profile settings are similar to the authentication settings used to set up a wireless network. For details and examples, see Configure VIA Client WLAN Profiles on page 691.
- **AAA profile:** This defines authentication settings for the WLAN users, including the role for unauthenticated users, and the different roles that should be assigned to users authenticated via 802.1x, MAC or SIP authentication. This profile includes references to:
  - MAC Authentication Profile
  - MAC Authentication Server Group
  - 802.1X Authentication Profile
  - 802.1X Authentication Server Group

- RADIUS Accounting Server Group

For details on configuring an AAA profile, see WLAN Authentication on page 421.

- **XML API server profile:** specifies the IP address of an external XML API server. For additional information, see Configuring the XML API Server on page 1024.
- **RFC 3576 server:** Specifies the IP address of a RFC 3576 RADIUS server. For additional information, see Configuring an RFC-3576 RADIUS Server on page 232.
- **MAC Authentication profile:** defines parameters for MAC address authentication, including upper- or lower-case MAC string, the diameter format in the string, and the maximum number of authentication failures before a user is blacklisted. For additional information, see Configuring the MAC Authentication Profile on page 250.
- **Captive Portal Authentication profile:** this profile directs clients to a web page that requires them to enter a username and password before being granted access to the network. This profile defines login wait times, the URLs for login and welcome pages, and manages the default user role for authenticated captive portal clients.

  You can also set the maximum number of authentication failures allowed per user before that user is blacklisted. This profile includes a reference to a Server group profile. For complete information on configuring a Captive portal authentication profile, refer to Captive Portal Authentication on page 299.
- **WISPr authentication profile:** WISPr authentication allows a "smart client" to authenticate on the network when they roam between Wireless Internet Service Providers, even if the wireless hotspot uses an ISP for which the client may not have an account. For more information on configuring WISPr authentication, see Configuring WISPr Authentication on page 289.
- **802.1X authentication profile:** defines default user roles for machine or 802.1X authentication, and parameters for 8021.X termination and failed authentication attempts. For a list of the basic parameters in the 802.1X authentication profile, refer to 802.1X Authentication on page 253
- **RADIUS server profile:** identifies the IP address of a RADIUS server and sets RADIUS server parameters such as authentication and accounting ports and the maximum allowed number of authentication retries. For a list of the parameters in the RADIUS profile, refer to Configuring a RADIUS Server on page 226
- **LDAP server profile:** defines an external LDAP authentication server that processes requests from the controller. This profile specifies the authentication and accounting ports used by the server, as well as administrator passwords, filters and keys for server access. For a list of the parameters in the LDAP profile, refer to Configuring an LDAP Server on page 232
- **TACACS server profile:** specifies the TCP port used by the server, the timeout period for a TACACS+ request, and the maximum number of allowed retries per user. For a list of the parameters in the TACACS profile, refer to Configuring a TACACS+ Server on page 234
- **Server group:** Tthis profile manages groups of servers for specific types of authentication. Server Groups identify individual authentication servers and let you create rules for clients based on attributes returned for the client by the server during authentication. For additional information on configuring server rules, see Configuring Server-Derivation Rules on page 242
- **VPN Authentication profile:** this profile identifies the default role for authenticated VPN clients and also references a server group. It also provides a separate VPN AAA authentication for a terminating remote AP (default-rap) and a campus AP (default-CAP). If you want to simultaneously deploy various combinations of a VPN client, RAP-psk, RAP-certs and CAP on the same controller, see Table 58.
- **Management Authentication profile:** enables or disables management authentication, and identifies the default role for authenticated management clients. This profile also references a server group. For more information on configuring a management authentication profile, see Management Authentication Profile Parameters on page 789.
- **Wired Authentication profile:** This profile merely references an AAA profile to be used for wired authentication. See Securing Wired Clients on page 865.

- **Stateful NTLM authentication Profile:** monitor the NTLM (NT LAN Manager) authentication messages between clients and an authentication server. If the client authenticates via an NTLM authentication server, the controller can recognize that the client has been authenticated and assign that client a specified user role. or details on configuring stateful authentication, see Stateful and WISPr Authentication on page 285.
- **Stateful Kerberos Authentication:** use stateful Kerberos authentication to configure a controller to monitor the Kerberos authentication messages between a client and a Windows authentication server. If the client successfully authenticates via an Kerberos authentication server, the controller can recognize that the client has been authenticated and assign that client a specified user role. For more information on stateful Kerberos authentication, see Configuring Stateful Kerberos Authentication on page 288.
- **Stateful 802.1X Authentication Profile:** enables or disables 802.1X authentication for clients on non-Dell APs, and defines the default role for those users once they are authenticated. This profile also references a server group to be used for authentication. For details on configuring stateful authentication, see Stateful and WISPr Authentication on page 285.

Mesh Profiles

You can provision Dell APs to operate as mesh points, mesh portals or remote mesh portals. The secure enterprise mesh environment routes network traffic between APs over wireless hops to join multiple Ethernet LANs or to extend wireless coverage. The Mesh profiles are:

- **Mesh high-throughput SSID profile**: enables or disables high-throughput (802.11n) features and 40 Mhz channel usage, and define values for aggregated MAC protocol data units (MDPUs) and Modulation and Coding Scheme (MCS) ranges. If none of the APs in your Mesh deployment are 802.11n-capable, you do not need to configure a mesh high-throughput SSID profile. For additional information on configuring this profile, see Creating and Editing Mesh High-Throughput SSID Profiles on page 556.
- **Mesh radio profile**: determines many of the settings used by mesh nodes to establish mesh links and the path to the mesh portal, including the maximum number of children a mesh node can accept, and transmit rates for the 802.11a and 802.11g radios. For additional information on configuring this profile, see Creating and Editing Mesh Radio Profiles on page 551.
- **Mesh cluster profile:** contains the mesh cluster name (MSSID), authentication methods, security credentials, and cluster priority. For additional information on configuring this profile, see Configuring Mesh Cluster Profiles on page 547.

QoS Profiles

The **QoS profiles** configure traffic management and VoIP functions.

- **WMM Traffic management profile:**the profile for Wi-Fi Multi-Media (WMM) traffic management prioritizes voice and video traffic above other data traffic . For additional information on configuring this profile, see Voice and Video on page 878.
- **Traffic management profile:** specifies the minimum percentage of available bandwidth to be allocated to a specific Virtual AP when there is congestion on the wireless network, and sets the interval between bandwidth usage reports. For additional information on configuring this profile, see Table 83.
- **VoIP call admission control profile:** Dell's Voice Call Admission Control limits the number of active voice calls per AP by load-balancing or ignoring excess call requests. This profile enables active load balancing and call admission controls, and sets limits for the numbers of simultaneous Session Initiated Protocol (SIP), SpectraLink Voice Priority (SVP), Cisco Skinny Client Control Protocol (SCCP), Vocera or New Office Environment (NOE) calls that can be handled by a single radio. For additional information on configuring this profile, see Scanning for VoIP-Aware ARM on page 918.

IDS Profiles

The IDS profiles manage settings for wireless intrusion protection (WIP) and The WLAN management system (WMS) on the controller that monitors wireless traffic to detect any new AP or wireless client station that tries to connect to the network. For details on IDS profile configuration settings, see Wireless Intrusion Prevention on page 452

HA Group profiles

This profile defines settings used by the high-availability:fast failover feature. For details, see Configuring High Availability on page 576

Other Profiles

The Controller profile and other profiles set the management password policy, define equipment OUIs and configure voice, video or VIA authentication settings.

- **VIA Authentication Profile**: define an authentication profile for the VIA feature.
- **VIA Connection Profile**: define authentication and connection settings profile for the VIA feature.
- **VIA Web Authentication**: define a VIA authentication profile to be used for Web authentication.
- **VIA Global Configuration**: select whether or not the controller should allow VIA SSL fallback.
- **Management Password Policy**: define a policy for creating management passwords.
- **Voip Logging**:enable voice logs by for a specific voice client based upon the client's MAC address. For details, see Advanced Voice Troubleshooting on page 933
- **SIP settings**:define a keep alive mechanism for the SIP sessions using the periodic session refresh request from the user agents. For details, see Understanding Extended Voice and Video Features on page 915
- **Dialplan Profile**: define SIP dial plans on the controller to provide outgoing PSTN calls.
- **Configure Real-Time Analysis**: enable real -time call quality analysis for voice calls. For details, see Understanding Extended Voice and Video Features on page 915
- **License Provisioning**: enable the centralized licensing feature. For details, see Centralized Licensing in a Multi-Controller Network on page 132
- **AirGroup AAA:** configure the AirGroup and ClearPass Policy Manager (CPPM) interface to allow an AirGroup controller and CPPM to exchange information about the owner, visibility, and status for each mobile device on the network. For details, see Configuring the AirGroup-CPPM Interface on page 965
- **CPPM IF-MAP:** use this feature in conjunction with ClearPass Policy Manager to send HTTP User Agent Strings and mDNS broadcast information to ClearPass so that it can make more accurate decisions about what types of devices are connecting to the network. For details, see ClearPass Profiling with IF-MAP on page 826.
- **Valid Equipment OUI Profile**: Set one or more Dell OUIs for the controller.
- **Upgrade**:configure the software upgrade feature that allows the master controller to automatically upgrade its associated local controllers by sending an image from a image server to one or more local controllers. For details, see Configuring Centralized Image Upgrades on page 792.

Profile Hierarchy

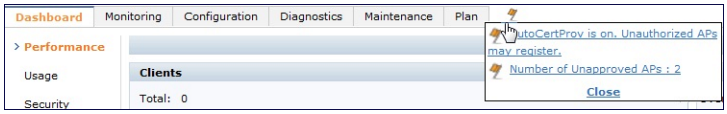
The ArubaOS WebUI includes several wizards that allow you to configure an AP, controller, WLAN, or License installation. You can also configure profiles using the WebUI Profile list or via the command line interface. Best practices is to configure the lowest-level settings first. For example, if you are defining a virtual AP profile, you should first define a session policy, then define your server group, then create an AAA profile that references the session policy and your server group.

The output of the **show profile-hierarchy** CLI command shows how profiles relate to each other, and how some higher-level profiles reference other lower-level profiles. The output of this command will vary, depending upon controller configuration and licenses.

Viewing Profile Errors

To view the list of profile errors using the CLI, use the **show profile-errors** command. The WebUI displays them with a *flag* icon next to the main horizontal menu (Figure 51). Click the flag to view the list of errors.

Figure 51 *Profile Errors*



Before you Deploy an AP

Before you install APs in a network environment, you must ensure that the APs are able to locate and connect to the controller. Specifically, you must configure firewall settings to allow APs to obtain software images and configuration settings from the controller, verify APs are able to locate the controller, and verify each AP is assigned a valid IP address when connected to the network. If you want to provision APs with more than one interface, you can also configure the USB settings and interface priority levels using an AP provisioning profile.

The following steps describe the basic pre-deployment tasks. Click any of the links for more information on these procedures.

1. Configure Firewall Settings
2. Enable Controller Discovery
3. Enable DHCP
4. (Optional) Define the AP Provisioning Profile
5. Define a virtual AP profile, and assign that profile to an AP group

Mesh AP Preconfiguration

Mesh APs require the following additional steps to define the mesh networking environment.

- Define and configure the mesh cluster profile.
- Define and configure the mesh radio profile

Remote AP Preconfiguration

Remote APs require the following additional step to identify valid APs in the remote AP whitelist.

- Create a Remote AP whitelist

Enable Controller Discovery

An AP can discover the IP address of the controller from a DNS server, from a DHCP server, or using the Aruba Discovery Protocol (ADP).

At boot time, the AP builds a list of controller IP addresses and then tries these addresses in order until it successfully reaches a controller . This list of IP addresses provides an enhanced redundancy scheme for controllers that are located in multiple data centers separated across Layer-3 networks. The AP constructs its list of controller addresses as follows:

- If the **master** provisioning parameter is set to a DNS name, that name is resolved and all resulting addresses are put on the list. If **master** is set to an IP address, that address is put on the list.
- If the **master** provisioning parameter is not set and a controller address was received in DHCP Option 43, that address is put on the list.
- If the **master** provisioning parameter is not set and no address was received via DHCP option 43, ADP is used to discover a controller address and that address is put on the list.
- Controller addresses derived from the **server-name** and **server-ip** provisioning parameters and the default controller name **aruba-master** are added to the list. Note that if a DNS name resolves to multiple addresses, all addresses are added to the list.

Controller Discovery using DNS

When using DNS, AP learns multiple IP addresses to associate with a controller. If the primary controller is unavailable or does not respond, the AP continues through the list of learned IP addresses until it establishes a connection with an available controller. This takes approximately 3.5 minutes per controller.



It is recommended you use a DNS server to provide APs with the IP address of the master controller because it involves minimal changes to the network and provides the greatest flexibility in the placement of APs.

APs are factory-configured to use the host name **aruba-master** for the master controller. For the DNS server to resolve this host name to the IP address of the master controller, you must configure an entry on the DNS server for the name **aruba-master**.

Controller Discovery using ADP

ADP is enabled by default on all Dell APs and controllers. With ADP, APs send out periodic multicast and broadcast queries to locate the master controller. ADP requires that all APs and controllers are connected to the same Layer-2 network. If the devices are on different networks, you must use a Layer-3 compatible discovery mechanism, such as DNS, DHCP, or IGMP forwarding.

To use ADP discovery:

1. Issue the command **show adp config** to verify that ADP and IGMP join options are enabled on the controller, If ADP is not enabled, you can reenable ADP using the command **adp discovery enable** and **adp igmp-join enable**.
2. If the APs are not in the same broadcast domain as the master controller, you enable multicast on the network (ADP multicast queries are sent to the IP multicast group address 239.0.82.11) for the controller to respond to the APs' queries. You also must make sure that all routers are configured to listen for Internet Group Management Protocol (IGMP) join requests from the controller and can route these multicast packets. C

Controller discovery using a DHCP Server

You can configure a DHCP server to provide the master controller's IP address. You must configure the DHCP server to send the controller's IP address using the DHCP vendor-specific attribute option 43. The APs identify themselves with a vendor class identifier set to **DellAP** in their DHCP requests. When the DHCP server responds to a request, it will send the controller's IP address as the value of option 43.

When using DHCP option 43, the AP accepts only one IP address. If the IP address of the controller provided by DHCP is not available, the AP can use the other IP addresses provisioned or learned by DNS to establish a connection. For more information on how to configure vendor-specific information on a DHCP server, see DHCP with Vendor-Specific Options on page 1055 or refer to the documentation included with your server.

Enable DHCP to Provide APs with IP Addresses

Each AP requires a unique IP address on a subnetwork that has connectivity to a controller. It is recommended you use the Dynamic Host Configuration Protocol (DHCP) to provide IP addresses for APs; the DHCP server can be an existing network server or an controller configured as a DHCP server.



If you do not enable DHCP, each AP must be manually configured with an IP address through the AP provisioning profile. For details, see .

You can use an existing DHCP server in the same subnetwork as the AP to provide the AP with its IP information. You can also configure a device in the same subnetwork to act as a relay agent for a DHCP server on a different subnetwork. (Refer to the vendor documentation for the DHCP Server or relay agent for information.)

If an AP is on the same subnetwork as the master controller, you can configure the controller as a DHCP server to assign an IP address to the AP. The controller must be the only DHCP server for this subnetwork.

In the WebUI

1. Navigate to the **Configuration > Network > IP > DHCP Server** window.
2. Select the **Enable DHCP Server** checkbox.
3. In the Pool Configuration section, click **Add**.
4. Enter information about the subnetwork for which IP addresses are to be assigned. Click **Done**.
5. If there are addresses that should not be assigned in the subnetwork:
   a. Click **Add** in the Excluded Address Range section.
   b. Enter the address range in the Add Excluded Address section.
   c. Click **Done**.
6. Click **Apply** at the bottom of the window.

In the CLI

```
(host)(config)# ip dhcp excluded-address ipaddripaddr2
(host)(config)# ip dhcp pool name
      default-router ipaddr
      dns-server ipaddr
      domain-name name
      network ipaddrmask
(host)(config)# service dhcp
```

AP Provisioning Profiles

AP provisioning profiles allow you to define a set of additional provisioning information for an AP, such as USB modem settings, PPPoE values, or configuration settings to provision an AP as a remote AP. When you create a provisioning profile, you can then apply that profile to an AP group and provision that entire group of campus or remote APs with the settings in that profile.

Defining an AP Provisioning Profile

By default, an AP group does not have a provisioning profile. Make sure that any provisioning profiles you create are complete and accurate before you assign that profile to an AP group. If a misconfigured provisioning profile is assigned to a group of APs, the APs in that group may be automatically provisioned with erroneous parameters and become lost.

1. Navigate to the **Configuration > Wireless > AP Installation > Provisioning** window.

2. Next, select the **Provisioning Profile** tab and enter a provisioning profile name in the text box (next to the Add button).
3. Click the **Add** button to add the profile name.
4. Select your new provisioning profile name from the list at the left.
5. (Optional) If you are provisioning a remote AP, select the **Remote-AP** checkbox.
6. Enter the IP address or the fully qualified domain name of the master controller in the **Master IP/FQDN** field.
7. If your AP will use Point-to-Point Protocol over Ethernet (PPPoE) to authenticate itself to a service provider, select the **PPPoE Parameters** checkbox and enter the following PPPoE values:
   - PPPoE User Name: Set the PPPoE User Name for this remote AP.
   - PPPoE Password: Enter and then confirm the PPPoE password for this remote AP.
   - PPPoE Service Name: Either an ISP name or a class of service configured on the PPPoE server.
8. (Optional) If you want to use this provisioning profile to provision APs with more than one interface, you must also configure the USB settings and priority levels for this profile. The configuration settings in this profile are described in Table 93.
9. Click **Apply** to save your settings.

Table 93: *AP Provisioning Profile parameters*

Parameter	Description
Remote-AP	Select this checkbox to provision the group of APs as remote APs.
Master IP/FQDN	The fully qualified domain name (FQDN) or IP address of the controller to which the AP is associated. **NOTE:** If you configure a master IP/FQDN setting in an AP's provisioning profile, this setting will override any LMS and backup LMS settings configured in an AP's AP system-profile. Leave the master IP/FQDN parameter blank if you want the AP to use the LMS or backup LMS values.
PPPOE User Name :	Point-to-Point Protocol over Ethernet (PPPoE) password for the AP.
PPPOE Password :	Point-to-Point Protocol over Ethernet (PPPoE) password for the AP.
PPPOE Service Name	Configures the PPPoE service name for the AP.
USB User Name	Configures the USB username for the AP.
USB Password :	A USB password, if provided by the cellular service provider.
USB Device Type	The USB device type.
USB Device Identifier	The USB device identifier.
USB Dial String	The dial string for the USB modem. This parameter only needs to be specified if the default string is not correct.
USB Initialization String	The initialization string for the USB modem. This parameter only needs to be specified if the default string is not correct.
USB TTY device data path	The TTY device path for the USB modem. This parameter only needs to be specified if the default path is not correct.
USB TTY device control path	The TTY device control path for the USB modem. This parameter only needs to be specified if the default path is not correct.

Parameter	Description
Link Priority Ethernet	Set the priority of the wired uplink. Each uplink type has an associated priority; wired ports having the highest priority by default.
Link Priority Cellular	Set the priority of the cellular uplink. By default, the cellular uplink is a lower priority than the wired uplink; making the wired link the primary link and the cellular link the secondary or backup link. Configuring the cellular link with a higher priority than your wired link priority will set your cellular link as the primary controller link.
Cellular modem network preference	The cellular modem network preference setting allows you to select how the modem should operate. • **auto** (default): In this mode, the modem firmware will control the cellular network service selection; so the cellular network service failover and fallback is not interrupted by the remote AP (RAP). • **3g_only**: Locks the modem to operate only in 3G. • **4g_only**: Locks the modem to operate only in 4G. • **advanced**: The RAP controls the cellular network service selection based on the Received Signal Strength Indication (RSSI) threshold-based approach. Initially the modem is set to the default auto mode. This allows the modem firmware to select the available network. The RAP determines the RSSI value for the available network type (for example 4G), checks whether the RSSI is within required range, and if so, connects to that network. If the RSSI for the modem's selected network is not within the required range, the RAP will then check the RSSI limit of an alternate network (for example, 3G), and reconnect to that alternate network. The RAP will repeat the above steps each time it tries to connect using a 4G multimode modem in this mode.
Username of AP so that AP can authenticate to 802.1x using PEAP	Configure the AP username.
Password of AP so that AP can authenticate to 802.1x using PEAP	Configure the AP password.
Uplink VLAN	If you configure an uplink VLAN on an AP connected to a port in trunk mode, the AP sends and receives frames tagged with this VLAN on its Ethernet uplink. By default, an AP has an uplink vlan of 0, which disables this feature. If an AP is provisioned with an uplink VLAN, it must be connected to a trunk mode port or the AP's frames will be dropped. 0 (disabled) to 4095 0
USB power mode	Set the USB power mode to control the power to the USB port.

Assigning Provisioning Profiles

Once you have defined a provisioning profile, you must assign that profile to an AP group.

1. Navigate to the **Configuration>AP configuration** window and select the **AP group** tab.
2. Click the **Edit** button by the name of the AP group to which you want to assign the provisioning profile.
3. In the profiles list, expand the **AP** menu, and select **Provisioning Profile.** The Profile Details window appears.
4. Click the **Provisioning Profile** drop-down list and select the name of the provisioning profile you want to assign to this AP group.
5. Click **Apply**.

If you are provisioning remote APs, you must also add the remote APs to the RAP whitelist. For details, see Remote Access Points on page 629.

Configuring Installed APs

APs and AMs are designed to require only minimal setup to make them operational in an user-centric network. Once APs have established communication with the controller, you can apply advanced configuration to individual APs or groups of APs in the network using the WebUI on the controller.

You can either connect the AP directly to a port on the controller, or connect the AP to another switch or router that has layer-2 or layer-3 connectivity to the controller. If the Ethernet port on the controller is an 802.3af Power over Ethernet (PoE) port, the AP automatically uses it to power up. If a PoE port is not available, you must get an AC adapter for the AP. For more information, see the Installation Guide for the specific AP.

If you are configuring a new AP that has never been provisioned before, you must first connect the AP to the controller according the instructions included with that AP. If you are reprovisioning or reconfiguring existing active APs, this step is not necessary, as the APs are already communicating with the controller.

This section describes the procedure to configure a installed AP with the basic settings it requires to become operational on the network. You can configure an AP using the AP wizard, the provisioning profile in the WebUI, or the controller command-line interface. using The individual configuration steps vary, depending upon whether the AP is deployed as a campus AP, remote AP (RAP) or a mesh AP.

Configuring an AP using the Provisioning Wizard

The easiest way to provision any AP is to use the AP Wizard in the controller WebUI. This wizard will walk you through the specific steps required to provision a campus, remote or Mesh AP. The Wizard includes a help tab that further describes each of the configuration tasks for that deployment type.

To access the AP wizard to provision a AP:

1. Select Configuration>Wizards>AP Wizard. The **Specify Deployment Scenario** window appears.
2. Select the deployment for the new AP, then click **Next** to continue to the next window in the Wizard.
3. Continue working your way through the Wizard to complete the provisioning process.

Configuring a AP using the WebUI

The following basic steps configure a campus AP on a LAN.



Remote APs and mesh APs require additional configuration steps not required for campus APs. For more information, see Configuring a Remote AP and

1. Navigate to the **Configuration > Wireless > AP Installation > Provisioning** window.
2. Click the checkbox by the AP you want to provision, then click **Provision**. The Provisioning window opens.
3. In the **AP Parameters** section, click the **AP Group** drop-down list and select the AP group to which this AP should be assigned. The AP group must have at least one virtual AP.
4. (Optional) Some AP models support an external antenna in addition to their internal antenna. If the AP you are provisioning supports an external antenna, the Provisioning window displays an additional **Antenna Parameters** section. If you want to use an External antenna for the remote AP you are provisioning, select **External Antenna** and define settings for that antenna. Otherwise, the remote AP will use its internal antenna by default.
5. If your AP will use Point-to-Point Protocol over Ethernet (PPPoE) to authenticate itself to a service provider, select the **PPPoE Parameters** checkbox and enter the following PPPoE values:
   - PPPoE User Name: Set the PPPoE User Name for this remote AP.

- PPPoE Password: Enter and then confirm the PPPoE password for this remote AP.
- PPPoE Service Name: Either an ISP name or a class of service configured on the PPPoE server.

6. (Optional) To allow the remote AP to use PEAP to authenticate to 802.1X networks, enter a user name and password in the 802.1X Parameter using PEAP section.
7. In the **Master Discovery** section, define how the AP should identify its WLAN controller. For more information on the different controller discovery methods, see Enable Controller Discovery on page 495.
8. (Optional) Define the uplink VLAN. If you configure an uplink VLAN on an AP connected to a port in trunk mode, the AP sends and receives frames tagged with this VLAN on its Ethernet uplink. To define the uplink VLAN, entering a VLAN ID from 1-4095 (inclusive) in the **IP Settings** section of the **Provisioning** window,
9. Define how the AP should obtain its IP address. If you have configured an DHCP server to allow APs to get addresses using DHCP, select **Obtain IP address using DHCP**. For more information on configuring a DHCP server, see Enable DHCP to Provide APs with IP Addresses on page 497. Otherwise, select **Use the Following IP address** and enter the appropriate values in the following fields:
   - **IP address**: IP address for the AP, in dotted-decimal format
   - **Subnet mask**: Subnet mask for the IP, in dotted-decimal format.
   - **Gateway IP address**: The IP address the AP uses to reach other networks.
   - **DNS IP address**: The IP address of the Domain Name Server.
   - **Domain name**: (optional) The default domain name.
10. (Optional) Access points can be configured in single-chain mode, allowing the radios of those APs to transmit and receive data using only legacy rates and single-stream HT and VHT rates on a single radio chain and single antenna or antenna interface. On APs with external antennas, this feature uses the external antenna interface labeled **A0** or **ANT0** (radio chain 0); the other (one or two) antenna interfaces are left unused. If you are provisioning an 802.11n-capable AP, select the **Enable for Radio-0** or **Enable for Radio-1** checkboxes in the **Single-Chain Mode** section to enable single-chain mode for the selected radio. AP radios in single-chain mode will transmit and receive data using only legacy rates and single-stream HT rates up to MCS 7. This feature is disabled by default.
11. (Optional) Define the AP name or SNMP location. The **AP list** section displays current information for an AP, and allows you to define additional parameters for your AP, such as AP Name, SNMP System Location.
12. Click **Apply and Reboot**. (Reprovisioning the AP causes it to automatically reboot).

Configuring a Remote AP

A remote AP (RAP) is recommended when the network between the AP and controller is an un-trusted/non-routable network, such as the Internet. Furthermore, a RAP supports an internal DHCP server, while a campus AP does not.

Remote Authentication

The two most common ways to provision an AP for remote authentication are certificate-based AP provisioning and provisioning using a pre-shared key. Although both options allow for a simple secure setup of your remote network, you should make sure that the procedure you select is supported by your controller, the AP model type and the end user's client software. If you must provision your APs using a pre-shared key, you need to know which controller models you have that do not support certificate-based provisioning.

- **Certificate based authentication** allows a controller to authenticate a AP using its certificates instead of a PSK. You can manually provision an individual AP with a full set of provisioning parameters, or simultaneously provision an entire group of APs by defining a provisioning profile which contains a smaller set of provisioning parameters that can be applied the entire AP group. When you manually provision an individual AP to use certificated-based authentication, you must connect that AP to the controller before you can define its provisioning settings.

- Use **Pre-Shared Key (PSK) authentication** to provision an individual remote AP or a group of remote APs using an Internet Key Exchange Pre-Shared Key (IKE PSK).

RAP Configuration

The steps to configure a remote AP using the WebUI are similar to the steps described in Configuring a AP using the WebUI , although some additional steps are required.

1. In the **Configuration > Wireless > AP Installation > Provisioning** window, select **Yes** for the **Remote AP** option.
2. In the **Remote IP Authentication Method** section, select either **Pre-shared key** or **certificate** authentication type. The Pre-shared key option requires you to perform the following additional steps:
   a. Enter and confirm the pre-shared key (IKE PSK).
   b. In the User credential assignment section, specify if you want to use a **Global User Name/password** or a **Per AP User Name/Password**.
      - If you use the Per AP User Names/Passwords option, each RAP is given its own user name and password.
      - If you use the Global User Name/Password option, all selected RAPs are given the same (shared) user name and password.
   c. Enter the user name, and enter and confirm the password. If you want the controller to automatically generate a user name and password, select **Use Automatic Generation**, then click **Generate** by the **User Name** and **Password** fields.
3. If you are provisioning remote AP models that support USB modems, you must complete the fields in the **USB settings** section. USB settings will not appear in the Provisioning window unless you are provisioning an AP that supports these features.

Configuring a Mesh AP

The steps to configure a remote AP using the WebUI are similar to the steps described in Configuring a AP using the WebUI , although some additional steps are required.

1. Define and configure the mesh cluster profile.
2. Define and configure the mesh radio profile
3. In the **AP list** section of the **Configuration > Wireless > AP Installation > Provisioning** window, select one of the following mesh for on the AP:
   - Mesh portal—The gateway between the wireless mesh network and the enterprise wired LAN.
   - Mesh point—APs that can provide traditional Dell WLAN services (such as client connectivity, intrusion detection system (IDS) capabilities, user roles association, LAN-to-LAN bridging, and Quality of Service (QoS) for LAN-to-mesh communication) to clients on one radio and perform mesh backhaul/network connectivity on the other radio. Mesh points can also provide LAN-to-LAN bridging through their Ethernet interfaces and provide WLAN services on the backhaul radio
   - Remote Mesh Portal: The Remote Mesh Portal feature allows you to configure a remote AP at a branch office to operate as a mesh portal for a mesh cluster.

For detailed provisioning guidelines, caveats, and instructions, see Secure Enterprise Mesh on page 533.

Verifying the Configuration

After the AP has been configured, navigate to **Monitoring>All Access Points** window and verify that the AP has an **up** status. The AP on your network *does not* appear in this table, it may have been classified as an inactive AP for any of the following reasons:

- The AP is configured with a missing or incorrect VLAN. (For example, the AP is configured to use a tunneled SSID of VLAN 2 but the controller doesn't have a VLAN 2.)
- The AP has an unknown AP group.
- The AP has a duplicate AP name.
- An AP with an external antenna is not provisioned with external antenna gain settings.
- Both radios on the AP are disabled.
- No virtual APs are defined on the AP.
- The AP has profile errors. For details, access the command-line interface and issue the command “show profile errors”.
- The GRE tunnel between the AP and the controller was blocked by a firewall after the AP became active.
- The AP is temporarily down while it is upgrading its software. The AP will become active again after upgrading.

Optional AP Configuration Settings

Once the AP has been installed and provisioned, you can use the WebUI or CLI to configure the optional AP settings described in the following sections:

- AP Installation Mode on page 503
- AP Name on page 504
- Spanning Tree on page 504
- RTLS Server on page 505
- AP Redundancy on page 506
- AP Maintenance Mode on page 507
- Energy Efficient Ethernet on page 507
- AP LEDs on page 508

AP Installation Mode

By default, all AP models initially ship with an indoor or outdoor installation mode. This means that APs with an indoor installation mode are normally placed in enclosed, protected environments and those with an outdoor installation mode are used in outdoor environments and exposed to harsh elements.

In most countries, there are different channels and power that are allowed for indoor and outdoor operation. You may want to change an AP’s installation mode from indoor to outdoor or vice versa.

Using the WebUI

To configure the installation mode for an AP, follow these steps:

1. Navigate to the **Configuration > Wireless> AP Installation** page. The list of discovered APs are displayed on this page.
2. Select the AP you want to change.
3. Click **Provision** to reveal the **Provisioning** page.

   Locate the **AP Installation Mode** section. By default, the **Default** mode is selected. This means that the AP installation type is based on the AP model.
4. Select the **Indoor** option to change the installation to Indoor mode. Select the **Outdoor** option to change the to Outdoor mode.
5. Click **Apply and Reboot** (at the bottom of the page).

Using the CLI

This example displays the AP installation mode options and sets the AP to indoor installation mode.

```
(host) (config) #provision-ap
(host) (AP provisioning) #installation ?
   default                  Decide by AP model
   indoor                   Indoor installation
   outdoor                  Outdoor installation
(host) (AP provisioning) #installation indoor
```

This example shows basic information details about the configuration of an AP named “MyAP.” The AP installation mode is indoor.

```
(host) #show ap details ap-name myAP

AP "MyAP" Basic Information
---------------------------
Item            Value
----            -----
AP IP Address   10.0.0.253
LMS IP Address  10.0.0.1
Group           default
Location Name   N/A
Status          Up; Mesh
Up time         9m:55s
Installation    indoor
```

AP Name

Each AP on the network should have a unique name. Display information about the APs on your network by executing the **show ap database long** command. The output will flag an AP that has a duplicate name (N flag).

Follow the steps below to rename an active AP on the network. If an AP with a duplicate name is no longer connected to your network, use the command **clear gap-db wired-mac** to clear the duplicate entry.

Using the WebU

1. Navigate to the **Configuration > Wireless> AP Installation** page. A list of discovered APs are on this page.
2. Select the AP you want to rename, and click **Provision**.
3. On the Provisioning page, scroll to the AP list at the bottom of the page and find the AP you want to rename.
4. In the AP Name field, enter the new unique name for the AP.
5. Click **Apply and Reboot**.

Using the CLI

Execute the following command (from enable mode) only on a master controller. Executing the command causes the AP to automatically reboot.

```
ap-rename {ap-name <name>|serial-num <number>|wired-mac <macaddr>} <new-name>
```

If an AP is recognized by the controller but is powered off or not connected to the network or controller when you execute the command, the request is queued until the AP is powered back on or reconnected.

Spanning Tree

The Spanning Tree Protocol (STP) can prevent loops in bridged Ethernet local area network. You can enable or disable the Spanning Tree parameter using the CLI and WebUI interfaces.

Using the WebUI

The following procedure configures the Spanning Tree parameter in AP System profile:

1. Navigate to the **Configuration > Advanced Services > All Profiles** page.
2. Under **AP > AP System** on the **Profiles** pane, select the profile name.
3. Under the **Basic** tab on the **Profile Details** pane, select the **Spanning Tree** checkbox.
4. Click **Apply**.

Using the CLI

STP is enabled only on wired ports of an AP. STP works only on downlink ports (eth1-<n>). The spanning Tree is supported in APs with 3 or more ports.

The following example enables spanning tree in default ap-system profile, using the CLI command:

```
(host) (config) #ap system-profile default
(host) (AP system profile "default") #spanning-tree
```

The following example displays the spanning tree information of an AP, using the CLI command:

```
(host) (config) #show ap debug spanning-tree ap-name <ap-name>
```

RTLS Server

The RTLS server configuration enables the AP to send RFID tag information to an RTLS server. Currently, when configuring the RTLS server under **ap system-profile**, the valid range of values for **station-message-frequency** was 5-3600 seconds. There are deployments that may require this to be configurable to as frequently as 1 per second. Starting with ArubaOS 6.4.2.0, you can set the **station-message-frequency** parameter in the 1-3600 seconds range. Setting the frequency to 1 means a report would be sent for every station every second. A value of 5 would mean that a report for any particular station would be sent at 5 second intervals.

In the WebUI

Use the following procedure to configure an RTLS server with station message frequency using the WebUI:

1. Navigate to the **Configuration > WIRELESS > AP Configuration** page.
2. Under the **AP Group** tab, click the desired profile.
3. Under the **Profiles** list, navigate to the **AP > AP system** profile menu.
4. Under the **Advanced** tab of the **Profile Details** section, configure the **RTLS Server configuration** parameters based on Table 94.

Table 94: *RTLS Server Configuration Parameters*

Parameter	Description
IPorDNS	IP address or the DNS of the RTLS server to which location reports are sent.
port	Port number on the server to which location reports are sent.
frequency	Indicates how often packets are sent to the server. Valid range is 1-3600 seconds.
key	Shared secret key.
includeUnassocSta	Indicates whether to include unassociated stations when sending station reports. Unassociated stations are stations that are not associated to any AP. Default value is disabled.

In the CLI

Use the following commands to configure an RTLS server with station message frequency using the CLI:

```
(host) (config) #ap system-profile default
(host) (AP system profile "default") #rtls-server ip-or-dns <IP or DNS of RTLS server> port
<port> key <key> station-message-frequency <1-3600>
```

Important Points to Remember

- Sending more frequent reports to the server can improve the accuracy of the location calculation.
- Configuring an AP to send reports more frequently adds additional load in terms of CPU usage.

AP Redundancy

In conjunction with the controller redundancy features described in Increasing Network Uptime Through Redundancy and VRRP on page 570 the information in this section describes redundancy for APs. Remote APs also offer redundancy solutions via a backup configuration, backup controller list, and remote AP failback. For more information relevant to remote APs, see Remote Access Points on page 629.

The AP failback feature allows an AP associated with the backup controller (backup LMS) to fail back to the primary controller (primary LMS) if it becomes available.

If configured, the AP monitors the primary controller by sending probes every 600 seconds by default. If the AP successfully contacts the primary controller for the entire hold-down period, it will fail back to the primary controller. If the AP is unsuccessful, the AP maintains its connection to the backup controller, restarts the LMS hold-down timer, and continues monitoring the primary controller.

The following example assumes:

- You have not configured the LMS or backup LMS IP addresses
- Default values unless otherwise noted.

Using the WebUI

Follow the procedure below to use the AP system profile to configure a redundant controller. For additional information on AP system profile settings, see .

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. Select either the AP Group or AP Specific tab. Click **Edit** for the AP group or AP name.
3. Under Profiles, select **AP** to display the AP profiles.
4. Select the AP system profile you want to modify.
5. Under Profile Details:
   a. At the **LMS IP** field, enter the primary controller IP address.
   b. At the **Backup LMS IP** field, enter the backup controller IP address.
   c. Click (select) **LMS Preemption**. This is disabled by default.
6. Click **Apply**.

Using the CLI

```
ap system-profile <profile>
   lms-ip <ipaddr>
   bkup-lms-ip <ipaddr>
   lms-preemption

ap-group <group>
   ap-system-profile <profile>

ap-name <name>
```

```
ap-system-profile <profile>
```

AP Maintenance Mode

You can configure APs to suppress traps and syslog messages related to those APs. Known as AP maintenance mode, this setting in the AP system profile is particularly useful when deploying, maintaining, or upgrading the network. If enabled, APs stop flooding unnecessary traps and syslog messages to network management systems or network operations centers during a deployment or scheduled maintenance. The controller still generates debug syslog messages if debug logging is enabled. After completing the network maintenance, disable AP maintenance mode to ensure all traps and syslog messages are sent. AP maintenance mode is disabled by default.

Using the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. Select either the AP Group or AP Specific tab. Click **Edit** for the AP group or AP name.
3. Under Profiles, select **AP** to display the AP profiles.
4. Select the AP system profile you want to modify.
5. Under Profile Details, do the following:
   - To enable AP maintenance mode, check (select) the **Maintenance Mode** checkbox.
   - To disable AP maintenance mode, clear (deselect) the **Maintenance Mode** checkbox.
6. Click **Apply**.

Using the CLI

To enable AP maintenance mode:

```
ap system-profile <profile>
   maintenance-mode
To disable AP maintenance mode:
ap system-profile <profile>
   no maintenance-mode
```

To view the maintenance mode status of APs, use the following commands:

```
show ap config {ap-group <name>|ap-name <name>|essid <name>}
show ap debug system-status {ap-name <name>|bssid <name>| ip-addr <ipaddr>}
```

On the local controller, you can also view maintenance mode status using the following commands:

```
show ap active {ap-name <name>|essid <name>|ip-addr <ipaddr>}
show ap database
show ap details {ap-name <name>|bssid <name>|ip-addr <ipaddr>}
```

Energy Efficient Ethernet

The W-AP130 Series support the 803.az Energy Efficient Ethernet (EEE) standard, which allows the APs to consume less power during periods of low data activity. This setting can be enabled for provisioned APs or AP groups through the Ethernet Link profile. If this feature is enabled for an APs group, any APs in the group that do not support 803.az will ignore this setting.

Using the WebUI

1. Navigate to the **Configuration > Advanced Services> All Profiles** page.
2. In the **Profiles** list, select **AP** to expand the AP profile menu.
3. Select **AP Ethernet Interface Link** profile. The list of existing Ethernet Link profiles appears in the **Profile Details** window. Select the Ethernet link profile you want to configure to support 803.az from this list, or create a new Ethernet link profile by entering a name for the new profile, then clicking **Add**.

4. The selected profile appears in the **Profile Details** window. The configuration parameters for the profile are described in Table 95.

Table 95: *Ethernet Interface Link Profile Parameters*

Parameter	Description
Speed	The speed of the Ethernet interface, either 10 Mbps, 100 Mbps, 1000 Mbps (1 Gbps), or auto-negotiated.
Duplex	The duplex mode of the Ethernet interface, either full, half, or auto-negotiated.
802.3az (EEE)	Select this checkbox to enable support for 802.1az Energy Efficient Ethernet. (for W-AP130 Series only).

5. Select the 803.az checkbox.
6. Click **Apply** to save your changes.

By default, AP wired port profiles reference the Default Ethernet interface link profile. If you created a new Ethernet interface link profile to support 803.az, use the procedure below to associate a AP wired port profile or Ethernet interface port configuration with the new Ethernet Interface link profile.

To associate a new Ethernet interface link profile with a wired port profile:

1. Navigate to the **Configuration > Advanced Services> All Profiles** page.
2. In the **Profiles** list, select **AP** to expand the **AP profile** menu.
3. Select **AP Wired Port Profile** to display a list of existing wired port profiles
4. Select the AP wired port profile you want to support 802.az. The Ethernet interface link profile currently associated with the port profile appears below the port profile in the **Profiles** list.
5. Click the Ethernet interface link profile currently associated with the AP wired port profile you want to modify. The settings for the Ethernet interface link profile appear in the **Profile Details** window.
6. Click the **Ethernet interface link profile** drop-down list at the top of the **Profile Details** window, and select a new Ethernet interface link profile.
7. Click **Apply** to save your changes.

Using the CLI

To enable support for 803.az EEE, access the command-line interface in config mode and issue the following command:

```
ap enet-link-profile <profile> dot3az
```

Associate a new Ethernet Interface link profile with an AP wired port profile using the following command:

ap wired-port-profile <profile>

```
enet-link-profile <profile>
```

AP LEDs

AP LEDs can be configured in two modes: **normal** and **off**. In normal mode, the AP LEDs will light as expected. When the mode is set to off, all of the LEDs on the affected APs are disabled.

Using the WebUI

An AP system profile's LED operating mode affects LEDS on all APs using that profile.



This option is available on the W-AP90 Series and W-AP105 access points.

1. Navigate to the **Configuration > Advanced Services> All Profiles** page.
2. Select the AP tab and then select the **AP system profiles** tab.
3. Select the AP system profile you want to modify.
4. Locate the **LED operating mode** parameter.
5. From the drop-down list, select **off**.
6. Click Apply.

Using the CLI

Use the **ap system-profile** command to disable LEDs for all APs using a particular system profile.

```
(host) (config)# ap system-profile <profile-name> led-mode {normal | off}
```

Use the **ap-leds** command to make the LEDs on a defined set of APs either blink or display in the currently configured LED operating mode. Note that if the LED operating mode defined in the AP's system profile is set to "off", then the **normal** parameter in the **ap-leds** command will disable the LEDs. If the LED operating mode in the AP system profile is set to "normal" then the **normal** parameter in this command will allow the LEDs light as usual.

```
(host) (config)# ap-leds {all | ap-group <ap-group> | ap-name <ap-name> | ip-addr <ip address>
| wired-mac <mac address>} {global blink|normal}|{local blink|normal}
```

RF Management

802.11a and 802.11g RF Management Profiles

The two 802.11a and 802.11g RF management profiles for an AP configure its 802.11a (5 Ghz) and 802.11b/g (2.4 GHz) radio settings. You can either use the "default" version of each profile, or create a new 802.11a or 802.11g profile using the procedures below. Each RF management radio profile includes a reference to an Adaptive Radio Management (ARM) profile. If you would like the ARM feature to dynamically select the best channel and transmission power for the radio, verify that the RF management profile references an active and enabled ARM profile. It can be useful to set the **Max Tx EIRP** parameter in the ARM profile to 127 (the maximum power level permissible) until it determines the signal-to-noise radio on the links. If ARM is active, the **Max Tx EIRP** can also be set to 127 to allow maximum power levels.

If you want to manually select a channel for each AP group, create separate 802.11a and 802.11g profiles for each AP group and assign a different transmission channel for each profile. For example, one AP group could have an 802.11a profile that uses channel 36 and an 802.11g profile that uses channel 11, and another AP group could have an 802.11a profile that uses channel 40 and an 802.11g profile that uses channel 9.

With the implementation of the high-throughput 802.11n standard, 40 MHz channels were added in addition to the existing 20 MHz channel options. Available 20 MHz and 40 MHz channels are dependent on the country code entered in the regulatory domain profile. The newer very high-throughput (VHT) 802.11ac standard introduces 80 Mhz channel options.



Changing the country code causes the valid channel lists to be reset to the defaults for the country.

The following channel configurations are available in ArubaOS:

- A 20 MHz channel assignment consists of a single 20 MHz channel. This channel assignment is valid for 802.11a/b/g and for 802.11n 20 MHz mode of operation.

- A 40 MHz channel assignment consists of two 20 MHz channels bonded together (a bonded pair). This channel assignment is valid for 802.11n 40 MHz mode of operation and is most often utilized on the 5 GHz frequency band.
- A 80 Mhz channel group for 5GHz radios. Only APs that support 802.11ac can be configured with 80 MHz channels.

  If high-throughput is disabled, a 40 MHz channel assignment can be configured, but only the primary channel assignment is used. The 20 MHz clients can also associate using this configuration, but only the primary channel is used.

APs initially start up with default **ack-timeout**, **cts-timeout** and **slot-time** values. When you modify the **maximum-distance** parameter in an rf dot11a radio profile or rf dot11g radio profile, new **ack-timeout**, **cts-timeout** and **slot-time** values may be derived, but those values are never less then the default values for an indoor AP.

Mesh radios on outdoor APs have additional constraints, as mesh links may need to span long distances. For mesh radios on outdoor APs, the effect of the default **maximum-distance** parameter on the **ack-timeout**, **cts-timeout** and **slot-time** values depends on whether the APs are configured as mesh portals or mesh points. This is because mesh portals use a default **maximum-distance** value of 16,050 meters, and mesh points use, by default, the maximum possible **maximum-distance** value.

The **maximum-distance** value should be set correctly to span the largest link distance in the mesh network so that when a mesh point gets the configuration from the network it will apply the correct **ack-timeout**, **cts-timeout** and **slot-time** values.The values derived from the **maximum-distance** setting depend on the band and whether 20Mhz/40MHz mode of operation is in use.

The following table indicates values for a range of distances:

Timeouts[usec]	--- 5GHz radio ---			--- 2.4GHz radio ---		
Distance[m]	Ack	CTS	Slot	Ack	CTS	Slot
0 (outdoor:16050m)	128	128	63	128	128	63
0 (indoor:600a,6450g)	25	25	9	64	48	9
200 (==default)	25	25	9	64	48	9
500	25	25	9	64	48	9
600	25	25	9	64	48	9
1050	28	28	13	64	48	31
5100	55	55	26	64	55	31
10050	88	88	43	88	88	43
15000	121	121	59	121	121	59
16050	128	128	63	128	128	63
58200(5G limit 20M)	409	409	203	-	-	-
52650(2.4G limit 20M)	-	-	-	372	372	185
27450(5G limit 40M)	204	204	101	-	-	-
24750(2.4G limit 40M)	-	-	-	186	186	92

VHT Support on W-AP200 Series, W-AP210 Series, W-AP220 Series, and W-AP270 Series Access Points

This feature enables Very High Throughput (VHT) rates on the 2.4 GHz band, providing 256-QAM modulation and encoding that allows for 600 Mbit/sec performance over 802.11n networks. Maximum data rates are increased on the 2.4 GHz band through the addition of VHT Modulation and Coding Scheme (MCS) values 8 and 9, which support the highly efficient modulation rates in 256-QAM. Starting with ArubaOS 6.4.2.0, VHT is supported on W-AP200 Series, W-AP210 Series, W-AP220 Series, and W-AP270 Series access points on both 20 MHz and 40 MHz channels.

Using the controller's CLI or WebUI, VHT MCS values 0-9 are enabled, overriding the existing high-throughput (HT) MCS values 0-7, which have a lower maximum data rate. However, this feature should be disabled if individual rate selection is required.

Managing 802.11a/802.11g Profiles Using the WebUI

Use the following procedures to define and manage 802.11a and 802.11g RF management profiles Using the WebUI.

Creating or Editing a Profile

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab.
   - If you selected **AP Group**, click the **Edit** button by the AP group for which you want to create or change an RF management profile.
   - If you selected **AP Specific**, click the **Edit** button by the AP for which you want to create or change an RF management profile.
2. In the Profiles list, expand the **RF Management** menu, then select either **802.11a radio profile** or **802.11g radio profile**.
3. To edit an existing 802.11a or 802.11g radio profile, select the desired profile from the **802.11a radio profile** or 802.11g radio profile drop-down list at the top of the **Profile Details** window. To create a new 802.11a or 802.11g profile, click the drop-down list at the top of the **Profile Details** window, select **NEW**, then enter a name for the new profile.

   The 802.11a and 802.11g profiles are divided into two tabs, **Basic** and **Advanced**. The **Basic** tab displays only those configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. If you change a setting on one tab then click and display the other tab without saving your configuration, that setting will revert to its previous value. The basic and advanced profile settings are described in Table 96.
4. Make the desired configuration changes, then click **Apply** to save your settings.

Table 96: *802.11a/802.11g RF Management Configuration Parameters*

Parameter	Description
Basic 802.11a/802.11g Settings – General	
Radio Enable	Enable transmissions on this radio band.
Mode	Access Point operating mode. Available options are: • **am-mode**: Air Monitor mode • **ap-mode**: Access Point mode • **spectrum-mode:** Spectrum Monitor mode The default settings is **ap-mode**.
High throughput enable (Radio)	Enable/Disable high-throughput (802.11n) features on the radio. This option is enabled by default.
Very high throughput enable (Radio)	Enable/Disable very high-throughput (802.11ac) features on the radio. This option is enabled by default. **NOTE:** This parameter is only available in the 802.11a radio profile.
Very high throughput rates enable (256-QAM)	Enable/Disable Very High Throughput (VHT) rate on 2.4 GHz band providing 256-QAM modulation and encoding that allows for 600 Mbit/sec performance over 802.11n networks. For more information, see VHT Support on W-AP200 Series, W-AP210 Series, W-AP220 Series, and W-AP270 Series Access Points. **NOTE:** This parameter is only available in the 802.11g radio profile.

Parameter	Description
Channel	Transmit channel for this radio. The available channels depend on the regulatory domain (country). This parameter includes the following channel number configuration options for 20 MHz, 40 MHz and 80 MHz modes: • **20**: Select this option to disable 40 MHz mode and 80 Mhz mode and activate 20 MHz mode for the entered channel. • **40**: Entering a channel number and selecting the **40** radio button in the WebUI selects a primary and secondary channel for 40 MHz mode. When you use this option, the number entered becomes the primary channel and the secondary channel is determined by increasing the primary channel number by 4. For example, if you entered 157 into the **Channel** field and selected the **above** option, radios using that profile would select 157 as the primary channel and 161 as the secondary channel. • **80**; Entering a channel number and selecting the 80 Mhz radio button selects a primary and secondary channel for 80 MHz mode. If you select the spectrum monitoring checkbox on this profile page, the AP will operate as a hybrid AP and scan the selected channel for spectrum analysis data.
Non-Wi-Fi Interference Immunity	Set a value for non-Wi-Fi Interference Immunity. The default setting for this parameter is level 2. When performance drops due to interference from non-802.11 interferers (such as DECT or Bluetooth devices), the level can be increased up to level 5 for improved performance. However, increasing the level makes the AP slightly "deaf" to its surroundings, causing the AP to lose a small amount of range. The levels for this parameter are: • Level 0: no ANI adaptation. • Level 1: noise immunity only. • Level 2: noise and spur immunity. • Level 3: level 2 and weak OFDM immunity. • Level 4: level 3 and FIR immunity. • Level 5: disable PHY reporting. **NOTE:** Only 802.11n-capable APs simultaneously support both the RX Sensitivity Tuning Based Channel Reuse feature and a level-3 to level-5 Noise Immunity setting. Do not raise the noise immunity default setting on APs that do not support 802.11n unless you first disable the Channel Reuse feature.
Spectrum Monitoring	Select this option to convert APs using this radio profile to a hybrid APs that will continue to serve clients as an Access Point, but will also scan and analyze spectrum analysis data for a single radio channel. For more details on hybrid APs, see Spectrum Analysis on page 700.
Advanced 802.11a/802.11g Settings	
Transmit EIRP	Maximum transmit EIRP in dBm from 0 to 51 in .5 dBm increments, or 127 for regulatory maximum. Transmit power may be further limited by regulatory domain constraints and AP capabilities.
Enable CSA	Channel Switch Announcements (CSAs), as defined by IEEE 802.11h, enable an AP to announce that it is switching to a new channel before it begins transmitting on that channel. This allows clients that support CSA to transition to the new channel with minimal downtime.
CSA Count	Number of channel switch announcements that must be sent prior to switching to a new channel. The default CSA count is 4 announcements.
Advertise 802.11d and 802.11h Capabilities	Enable the radio to advertise its 802.11d (Country Information) and 802.11h (Transmit Power Control) capabilities. This option is disabled by default.

Parameter	Description
Spectrum Load Balancing	The Spectrum Load Balancing feature helps optimize network resources by balancing clients across channels, regardless of whether the AP or the controller is responding to the wireless clients' probe requests. If enabled, the controller compares whether or not an AP has more clients than its neighboring APs on other channels. If an AP's client load is at or over a predetermined threshold as compared to its immediate neighbors, or if a neighboring Dell AP on another channel does not have any clients, load balancing will be enabled on that AP. This feature is disabled by default. For details, see Spectrum Load Balancing on page 449.
Beacon Period	Beacon Period for the AP in msec. The minimum value is 60 msec, and the default value is 100 msec.
Beacon Regulate	Enable this setting to introduce randomness in the beacon generation so that multiple APs on the same channel do not send beacons at the same time.
Advertised Regulatory Max EIRP	Work around a known issue on Cisco 7921G telephones by specifying a cap for a radio's maximum equivalent isotropic radiated power (EIRP). When you enable this parameter, even if the regulatory approved maximum for a given channel is higher than this EIRP cap, the AP radio using this profile will advertise only this capped maximum EIRP in its radio beacons. The supported range is 1-31dBm.
ARM/WIDS Override	If selected, this option disables Adaptive Radio Management (ARM) and Wireless IDS functions and slightly increases packet processing performance. If a radio is configured to operate in Air Monitor mode, then the ARM/WIDS functions are always enabled, regardless of whether or not this check box is selected.
Reduce Cell Size (Rx Sensitivity)	The cell size reduction feature allows you manage dense deployments and to increase overall system performance and capacity by shrinking an AP's receive coverage area, thereby minimizing co-channel interference and optimizing channel reuse. This value should only be changed if the network is experiencing performance issues. The possible range of values for this feature is 0-55 dB. The default 0 dB reduction allows the radio to retain its current default Rx sensitivity value. Values from 1 dB - 55 dB reduce the power level that the radio can hear by that amount. If you configure this feature to use a non-default value, you must also reduce the radio's transmission (Tx) power to match its new received (Rx) power level. Failure to match a device's Tx power level to its Rx power level can result in a configuration that allows the radio to send messages to a device that it cannot hear.
Management Frame Throttle Interval	Averaging interval for rate limiting management frames from this radio, in seconds. A management frame throttle interval of 0 seconds disables rate limiting.
Management Frame Throttle Limit	Maximum number of management frames that can come in from this radio in each throttle interval.
Maximum Distance	Maximum client distance, in meters. This value is used to derive ACK and CTS timeout times. A value of 0 specifies default settings for this parameter, where timeouts are only modified for outdoor mesh radios which use a distance of 16km. The upper limit for this parameter varies from 24km–58km, depending on the radio's band (a/g) and 20/40 MHz mode. Note that if you configure a value above the supported maximum, the maximum supported value will be used instead. Values below 600m will use default settings.

Parameter	Description
RX Sensitivity Tuning Based Channel Reuse	In some dense deployments, it is possible for APs to hear other APs on the same channel. This creates co-channel interference and reduces the overall utilization of the channel in a given area. Channel reuse enables dynamic control over the receive (Rx) sensitivity in order to improve spatial reuse of the channel. This feature is disabled by default. To enable this feature, click the **RX Sensitivity Tuning Based Channel Reuse** drop-down list and select either **static** or **dynamic**.To disable this feature, click the RX Sensitivity Tuning Based Channel Reuse drop-down list and select **disable**. For details on each of these modes, see Reusing Channels to Control RX Sensitivity Tuning on page 449. **NOTE:** Do not enable the Channel Reuse feature if Non-Wi-Fi Interference Immunity on page 512 is set to level 3 or higher. A level-3 to level-4 Noise Immunity setting is not compatible with the Channel Reuse feature. The channel reuse feature applies to non-DFS channels only. It is internally disabled for DFS channels and is does not affect DFS radar signature detection.
RX Sensitivity Threshold	RX sensitivity tuning based channel reuse threshold, in - dBm. If the Rx Sensitivity Tuning Based Channel reuse feature is set to static mode, this parameter manually sets the AP's Rx sensitivity threshold (in -dBm). The AP will filter out and ignore weak signals that are below the channel threshold signal strength. If the value for this parameter is set to zero, the feature will automatically determine an appropriate threshold.
Protection for 802.11b Clients	**(For 802.11g RF Management Profiles only)** Enable or disable protection for 802.11b clients. This parameter is enabled by default. Disabling this feature may improve performance if there are no 802.11b clients on the WLAN. WARNING: Disabling protection violates the 802.11 standard and may cause interoperability issues. If this feature is disabled on a WLAN with 802.11b clients, the 802.11b clients will not detect an 802.11g client talking and can potentially transmit at the same time, thus garbling both frames.
Associated Profiles	
ARM profile	Dell's proprietary Adaptive Radio Management (ARM) technology maximizes WLAN performance by dynamically and intelligently choosing the best 802.11 channel and transmit power for each Dell AP in its current RF environment. Every RF management profile references an ARM profile. If you specify an active and enabled ARM profile, you do not need to manually configure the **Channel** and **Transmit Power** parameters for this 802.11a or 802.11g profile. For details on referencing an ARM profile, see Assigning an ARM Profile on page 516. The Adaptive Radio Management (ARM) profile associated with this 802.11a or 802.11g radio profile appears beneath the 802.11a/802.11g radio profile name in the profiles list. To change the ARM profile associated with an 802.11a or 802.11g radio profile, select the associated ARM profile in the profiles list then click the drop-down list in the **Profile Details** section of the page to select a new profile.

Parameter	Description
High-throughput radio profile	A high-throughput profile manages 40 MHz tolerance settings, and controls whether or not APs using this profile will advertise intolerance of 40 MHz operation. (This option is disabled by default, allowing 40 MHz operation.) A high-throughput profile also determines whether an AP radio using the profile will stop using the 40 MHz channels surrounding APs or stations advertise 40 MHz intolerance. This option is enabled by default. For details on referencing a high-throughput radio profile, see Assigning a High-throughput Profile on page 515. The high-throughput radio profile associated with this 802.11a or 802.11g radio profile appears beneath the 802.11a/802.11g radio profile name in the profiles list. To change the high-throughput radio profile associated with an 802.11a or 802.11g radio profile, select the associated high-throughput radio profile in the profiles list then click the drop-down list in the **Profile Details** section of the page to select a new profile.
Spectrum Monitoring Profile	The spectrum monitoring profile defines the spectrum band and device ageout times used by a spectrum monitor radio. The spectrum monitoring profile associated with this 802.11a or 802.11g radio profile appears beneath the 802.11a/802.11g radio profile name in the profiles list. To change the spectrum monitoring profile associated with an 802.11a or 802.11g radio profile, select the associated spectrum monitoring profile in the profiles list then click the drop-down list in the **Profile Details** section of the page to select a new profile.
AM Scanning Profile	The AM scanning profile associated with this 802.11a or 802.11g radio profile appears beneath the 802.11a/802.11g radio profile name in the profiles list. To change the AM scanning profile associated with an 802.11a or 802.11g radio profile, select the associated AM scanning profile in the profiles list then click the drop-down list in the **Profile Details** section of the page to select a new profile.

Assigning an 802.11a/802.11g Profile to an AP or AP Group

Use the following procedure to assign an 802.11a or 802.11g RF management profile to an AP Group or individual AP using the WebUI.

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab.
   - If you selected **AP Group**, click the **Edit** button by the AP group name to which you want to assign a new 802.11a or 802.11g RF management profile.
   - If you selected **AP Specific**, click the **Edit** button by the AP to which you want to assign a new 802.11a or 802.11g RF management profile
2. Under the **Profiles** list, expand the **RF management** menu, then select either **802.11a radio profile** or e**802.11g radio profile**.
3. To select a 802.11a radio profile for an AP or AP group, click the **802.11a radio profile** drop-down list in the **Profile Details** window pane and select the desired profile from the list.
   -or-
   To select a 802.11g radio profile for an AP or AP group, click the **802.11g radio profile** drop-down list in the **Profile Details** window pane and select the desired profile from the list.
4. Click **Apply**. The profile name appears in the Profile list with your configured settings. If you configure this for the AP group, this profile also becomes the selected 802.11a or 802.11g RF management profile used by the mesh portal for your mesh network.

Assigning a High-throughput Profile

Each 802.11a or 802.11g RF management radio profile references a high-throughput profile that manages the AP group's 40Mhz tolerance settings. By default, an 802.11a profile references a high-throughput profile

named default-a and an 802.11g profile references a high-throughput profile named default-g. If you do not want to use these default profiles, use the procedure below to reference a different high-throughput profile for your 802.11a or 802.11g RF management profiles.

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab.
   - If you selected **AP Group**, click the **Edit** button by the AP group name to which you want to assign a new high-throughput profile.
   - If you selected **AP Specific**, click the **Edit** button by the AP which you want to assign a new high-throughput profile.
2. In the **Profiles** list, expand the **RF Management** menu, then select either **802.11a radio profile** or **802.11g radioprofile**.
3. Select**High-throughput radio profile**.The **Profile Details** pane appears and displays information for the currently referenced high-throughput profile. Use this window pane to select a different high-throughput profile, or to create an entirely new high-throughput profile for that 802.11a or 802.11g radio.
   - To reference a different high-throughput profile, click the **High-throughput Radio Profile** drop-down list and select a new profile name from the list. Click **Apply** to save your changes.
   - To create a new high-throughput profile, click the **High-throughput Radio Profile** drop-down list and select **NEW**.

     a. Enter a name for the new high-throughput profile.

     b. *(Optional)* Select **40 MHz intolerance** if you want to enable 40 MHz intolerance. This parameter controls whether or not APs using this high-throughput profile will advertise intolerance of 40 MHz operation. By default, this option is disabled and 40 MHz operation is allowed.

     d. *(Optional)* Select **honor40 MHz intolerance** to allow a radio using this profile to stop using the 40 MHz channels if the 40 MHz intolerance indication is received from another AP or station. This option is enabled by default.

     d. Click **Apply** to save your settings.
4. The high-throughput profile appears in the **Profile** list with your configured settings.

Assigning an ARM Profile

By default, an 802.11a or 802.11g profile references an ARM profile named **default**. Most network administrators will find that this one default ARM profile is sufficient to manage all the Dell APs on their WLAN. If, however, you do not want to use this default ARM profile, use the procedure below to reference a different ARM profile for your 802.11a or 802.11g RF management profiles.

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab.
   - If you selected **AP Group**, click the **Edit** button by the AP group name to which you want to assign a new ARM profile.
   - If you selected **AP Specific**, click the **Edit** button by the AP which you want to assign a new ARM profile.
2. Under the Profiles list, expand the **RF Management** menu.
3. To reference an ARM profile for a 802.11a radio profile, expand the **802.11a radio** profile menu.

   -or-

   To reference an ARM profile for a 802.11g radio profile, expand the **802.11g radio** profile menu.
4. The **Profile Details** pane appears and displays information for the currently referenced ARM profile. You can now select a different profile, or create an entirely new ARM profile for that 802.11a or 802.11g radio.
   - To reference a different ARM profile, click the **Adaptive Radio Management (ARM)** Profile drop-down list and select a new profile name from the list. Click Apply to save your changes.

- To create a new ARM profile, click the **Adaptive Radio Management (ARM)** Profile drop-down list and select **NEW**.

  a. Enter a name for your new ARM profile.

  b. (Optional) If you are not configuring ARM for a mesh node, select **40 MHz intolerance** if you want to enable 40 MHz intolerance. This parameter controls whether or not APs using this high-throughput profile will advertise intolerance of 40 MHz operation. By default, this option is disabled and 40 MHz operation is allowed.

  c. (Optional)If you are not configuring ARM for a mesh node, select **honor 40 MHz intolerance** to allow a radio using this profile to stop using the 40 MHz channels if the 40 MHz intolerance indication is received from another AP or station. This option is enabled by default.

5. Click**Apply** to save your settings.

The ARM profile name appears in the Profile list with your configured settings. If you configured this profile for the AP group, this ARM profile becomes part of the selected 802.11a or 802.11g RF management profile used by the mesh portal for your mesh network.

Deleting a Profile

You can delete an 802.11a or 802.11g radio profile only if no APs or AP groups are associated with that profile. To delete a 802.11a or 802.11g radio profile using the WebUI.

1. Navigate to the **Configuration > Advanced Services> All Profiles** window.
2. Expand the **RF Management** menu, then select **802.11a radio profile** or **802.11g radio profile**. A list of profiles of the specified type appears in the **Profile Details** window pane.
3. Click the **Delete** button by the name of the profile you want to delete.

Managing 802.11a/802.11g Profiles Using the CLI

You must be in config mode to create, modify or delete a 802.11a or 802.11g RF management radio profile using the CLI. Specify an existing mesh profile with the <profile-name> parameter to modify an existing profile, or enter a new name to create an entirely new profile.

Creating or Modifying a Profile

Configuration details and any default values for each of these parameters are described in Table 96. This CLI command also allows you to reference an ARM profile and high-throughput radio profile for the 802.11a or 802.11g radio. If you do not specify a parameter for a new profile, that profile uses the default value for that parameter. Put the `no` option before any parameter to remove the current value for that parameter and return it to its default setting. Enter `exit` to leave the 802.11a or 802.11g profile mode.

```
rf dot11a-radio-profile|dot11g-radio-profile <profile-name>
   am-scan-profile
   arm-profile
   beacon-period
   beacon-regulate
   cap-reg-eirp
   channel <num|num+|num->
   channel-reuse
   channel-reuse-threshold
   clone
   csa
   csa-count
   disable-arm-wids-function
   dot11b-protection (for 802.11g radio profiles only)
   dot11h
   high-throughput-enable
   ht-radio-profile
```

```
interference-immunity
maximum-distance
mgmt-frame-throttle-interval
mgmt-frame-throttle-limit
mode {ap-mode|am-mode|spectrum-mode}
no
radio-enable
slb-mode
slb-threshold
slb-update-interval
spectrum-load-bal-domain
spectrum-load-balancing
spectrum-monitoring
spectrum-profile
tpc-power
tx-power
```

You can also create a new 802.11a or 802.11g RF management profile by copying the settings of an existing profile using the clone parameter. Using the `clone` command to create a new profile makes it easier to keep constant attributes in common within multiple profiles.

```
rf dot11a-radio-profile <profile-name>   clone <source-profile-name>
rf dot11g-radio-profile <profile-name>   clone <source-profile-name>
```

Viewing RF Management Settings

To view a complete list of 802.11a or 802.11g RF management profiles and their status:

```
show rf dot11a-radio-profile|dot11g-radio-profile
```

To view the settings of a specific RF management profile:

```
show rf dot11a-radio-profile|dot11g-radio-profile <profile-name>
```

Assigning a 802.11a/802.11g Profile

To assign an 802.11a or 802.11g RF management profile to an AP group:

```
ap-group <group> dot11a-radio-profile <profile-name>
```

-or-

```
ap-group <group> dot11g-radio-profile <profile-name>
```

To assign an 802.11a or 802.11g RF management profile to an individual AP:

```
ap-name <name> dot11a-radio-profile <profile-name>
```

-or-

```
ap-name <name> dot11g-radio-profile <profile-name>
```

Deleting a Profile

If no AP or AP group is using an RF management profile, you can delete that profile using the **no** parameter:

```
no rf dot11a-radio-profile <profile-name>
```

RF Optimization

Each AP includes an RF Optimization profile that allows you to configure settings for detecting interference. The controller can detect interference near a wireless client station or AP is based on an increase in the frame retry rate or frame receive error rate.

Using the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab.

- If you selected the **AP Group** tab, click the **Edit** button by the AP group name for which you want to configure the RF Optimization profile.
- If you selected the **AP Specific** tab, click the **Edit** button by the AP for which you want to create the RF Optimization profile.

2. Expand the **RF Management** menu, then expand the **RF Optimization Profile** menu.
3. Select the profile you want to edit from the **Profile Details** window pane.

   or

   Enter a new RF Optimization profile name in the field at the bottom of the **Profile Details** window, then click **Add**. Next, select that profile name from the profile list to edit its parameters.
4. Configure your RF Optimization radio settings. Table 97 describes the parameters. Click **Apply** to save your settings.

Table 97: *RF Optimization Profile Parameters*

Parameter	Description
Station Handoff Assist	Allows the controller to force a client off an AP when the RSSI drops below a defined minimum threshold. Default: Disabled
RSSI Falloff Wait Time	Time, in seconds, to wait with decreasing RSSI before a de-authorization message is sent to the client. Maximum value: 8 seconds Default value: 0 seconds
Low RSSI Threshold	Minimum RSSI above which de-authorization messages should never be sent.
RSSI Check Frequency	Interval, in seconds, to sample RSSI.

Using the CLI

Use the following command to configure RF Optimization profiles. The parameters described in Table 97.

```
rf optimization-profile <profile>
clone <profile>
handoff-assist
low-rssi-threshold <number>
no ...
rssi-check-frequency <number>
rssi-falloff-wait-time <seconds>
```

RF Event Configuration

An AP's event threshold profile configures Received Signal Strength Indication (RSSI) metrics, including high and low watermarks for frame error rates and frame retry rates. When certain RF parameters are exceeded, these events can signal excessive load on the network, excessive interference, or faulty equipment.



This profile and many of the detection parameters are disabled (value is 0) by default.

The following procedure details the steps to configure RF Event parameters.

Using the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab.

- If you selected the **AP Group** tab, click the **Edit** button by the AP group name for which you want to configure the RF Event profile.
- If you selected the **AP Specific** tab, click the **Edit** button by the AP for which you want to create the RF Event profile.

2. In the Profiles list, expand the **RF Management** menu, then expand the **RF Event Profile** menu.
3. To edit an existing RF Event profile, select the profile you want to edit from the **Profile Details** window pane.

   -or-
4. To create a new profile, enter a new RF Event profile name in the field at the bottom of the **Profile Details** window, then click **Add**. Next, select that profile name from the profile list to edit its parameters.
5. Configure your settings as detailed in Table 98 and click **Apply** to save your settings.

Table 98: *RF Event Thresholds Profile Parameters*

Parameter	Description
Detect Frame Rate Anomalies	Enable or disables detection of frame rate anomalies. This feature is disabled by default.
Bandwidth Rate High Watermark	If bandwidth in an AP exceeds this value, a bandwidth exceeded condition exists. The value represents the percentage of maximum for a given radio. (For 802.11b, the maximum bandwidth is 7 Mbps. For 802.11 a and g, the maximum is 30 Mbps.) The recommended value is 85%.
Bandwidth Rate Low Watermark	After a bandwidth exceeded condition exists, the condition persists until bandwidth drops below this value. The recommended value is 70%.
Frame Error Rate High Watermark	If the frame error rate (as a percentage of total frames in an AP) exceeds this value, a frame error rate exceeded condition exists. The recommended value is 16%.
Frame Error Rate Low Watermark	After a frame error rate exceeded condition exists, the condition persists until the frame error rate drops below this value. The recommended value is 8%.
Frame Fragmentation Rate High Watermark	If the frame fragmentation rate (as a percentage of total frames in an AP) exceeds this value, a frame fragmentation rate exceeded condition exists. The recommended value is 16%.
Frame Fragmentation Rate Low Watermark	After a frame fragmentation rate exceeded condition exists, the condition persists until the frame fragmentation rate drops below this value. The recommended value is 8%
Frame Low Speed Rate High Watermark	If the rate of low-speed frames (as a percentage of total frames in an AP) exceeds this value, a low-speed rate exceeded condition exists. This could indicate a coverage hole. The recommended value is 16%.
Frame Low Speed Rate Low Watermark	After a low-speed rate exceeded condition exists, the condition persists until the percentage of low-speed frames drops below this value. The recommended value is 8%.
Frame Non Unicast Rate High Watermark	If the non-unicast rate (as a percentage of total frames in an AP) exceeds this value, a non-unicast rate exceeded condition exists. This value depends upon the applications used on the network.
Frame Non Unicast Rate Low Watermark	After a non-unicast rate exceeded condition exists, the condition persists until the non-unicast rate drops below this value.

Parameter	Description
Frame Receive Error Rate High Watermark	If the frame receive error rate (as a percentage of total frames in an AP) exceeds this value, a frame receive error rate exceeded condition exists. The recommended value is 16%
Frame Receive Error Rate Low Watermark	After a frame receive error rate exceeded condition exists, the condition persists until the frame receive error rate drops below this value. The recommended value is 8%.
Frame Retry Rate High Watermark	If the frame retry rate (as a percentage of total frames in an AP) exceeds this value, a frame retry rate exceeded condition exists. The recommended value is 16%.
Frame Retry Rate Low Watermark	After a frame retry rate exceeded condition exists, the condition persists until the frame retry rate drops below this value. The recommended value is 8%.

Using the CLI

Use the following command to configure RF event profiles. The available parameters for this profile are detailed in Table 98.

```
rf event-thresholds-profile <profile>
bwr-high-wm <percent>
bwr-low-wm <percent>
clone <profile>
detect-frame-rate-anomalies
fer-high-wm <percent>
fer-low-wm <percent>
ffr-high-wm <percent>
ffr-low-wm <percent>
flsr-high-wm <percent>
flsr-low-wm <percent>
fnur-high-wm <percent>
fnur-low-wm <percent>
frer-high-wm <percent>
frer-low-wm <percent>
frr-high-wm <percent>
frr-low-wm <percent>
```

Optimizing APs Over Low-Speed Links

Depending on your deployment scenario, you may have APs or remote APs that connect to a controller located across low-speed (less than 1 Mbps capacity) or high-latency (greater than 100 ms) links.

With low-speed links, if heartbeat or keep alive packets are not received between the AP and controller during the defined interval, APs may reboot causing clients to re-associate. You can adjust the bootstrap threshold and prioritize AP heartbeats to optimize these types of links. In addition, high bandwidth applications may saturate low-speed links. For example, if you have tunnel-mode SSIDs, use them with low-bandwidth applications such as barcode scanning, small database lookups, and Telnet to avoid saturating the link. If you have traffic that will remain local, deploying remote APs and configuring SSIDs as bridge-mode SSIDs can also prevent link saturation.

With high-latency links, consider the amount and type of client devices accessing the links. Dell APs locally process 802.11 probe-requests and probe-responses, but the 802.11 association process requires interaction with the controller.

When deploying APs across low-speed or high-latency links, The following best practices are recommended:

- Connect APs and controllers over a link with a capacity of 1 Mbps or greater.
- Maintain a minimum link speed of 64 Kbps per AP and per bridge-mode SSID. This is the minimum speed required for downloading software images.
- Adjust the bootstrap threshold to 30 if the network experiences packet loss. This makes the AP recover more slowly in the event of a failure, but it will be more tolerant to heartbeat packet loss.
- Prioritize AP heartbeats to prevent losing connectivity with the controller.
- If possible, reduce the number of tunnel-mode SSIDs. Each SSID creates a tunnel to the controller with its own tunnel keep alive traffic.
- If most of the data traffic will remain local to the site, deploy remote APs in bridging mode. For more information about remote APs, see Access Points (APs) on page 485.
- If high-latency links such as transoceanic or satellite links are used in the network, deploy a controller geographically close to the APs.
- If high-latency causes association issues with certain handheld devices or barcode scanners, check the manufacturer of the device for recent firmware and driver updates.

Configuring the Bootstrap Threshold

To configure the bootstrap threshold using the WebUI:

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. Select either the **AP Group** or **AP Specific** tab. Click **Edit** by the AP group or AP name.

   The AP system profile configuration settings are divided into two tabs, **Basic** and **Advanced**. The **Basic** tab displays only those configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. If you change a setting on one tab then click and display the other tab without saving your configuration, that setting will revert to its previous value. Both basic and advanced settings are described in Table 99.
3. Under Profiles, select **AP**, then **AP system profile**. The profile appears the **Profile Details** window.
4. In the **Bootstrap threshold** field, enter 30.
5. Click **Apply**.

Table 99: *AP System Profile Configuration*

Parameter	Description
Basic AP System Profile Settings–General	
RF Band	For APs that support both 802.11a and 802.11b/g RF bands, specify the RF band in which the AP should operate: • g = 2.4 GHz • a = 5 GHz
RF Band for AM Mode scanning	For Air Monitors that support both 802.11a and 802.11b/g RF bands, specify the RF band which the AM should scan: • a = 5 GHz • all = both radio bands • g = 2.4 GHz
Native VLAN ID	Native VLAN for bridge mode virtual APs (frames on the native VLAN are not tagged with 802.1q tags).
Session ACL	Session ACL configured with the ip access-list session command.

Parameter	Description
	NOTE: This parameter requires the PEFNG license.
Corporate DNS Domain	Name of domain that is resolved by corporate DNS servers. Use this parameter when configuring split-tunnel forwarding.
SNMP sysContact	SNMP system contact information.
LED operating mode	The operating mode for the 802.11n-capable AP LEDs.
Basic AP System Profile Settings–LMS	
SAP MTU	Maximum Transmission Unit, in bytes, on the wired link for the AP.
LMS IP	In multi-controller networks, this parameter specifies the IP address of the local management switch (LMS)—the Dell controller—which is responsible for terminating user traffic from the APs, and processing and forwarding the traffic to the wired network. This can be the IP address of the local or master controller. When using redundant controllers as the LMS, set this parameter to be the VRRP IP address to ensure that APs always have an active IP address with which to terminate sessions. **NOTE:** If the LMS-IP is blank, the access point will remain on the controller that it finds using methods like DNS or DHCP. If an IP address is configured for the LMS IP parameter, the AP will be immediately redirected to the controller at that address.
Backup LMS IP	In multi-controller networks, specifies the IP address of a *backup* to the IP address specified with the lms-ip parameter.
LMS IPv6	In multi-controller ipv6 networks, specifies the IPv6 address of the local management switch (LMS)—the controller—which is responsible for terminating user traffic from the APs, and processing and forwarding the traffic to the wired network. This can be the IP address of the local or master controller. When using redundant controllers as the LMS, set this parameter to be the VRRP IP address to ensure that APs always have an active IP address with which to terminate sessions.
Backup LMS IPv6	In multi-controller ipv6 networks, specifies the IPv6 address of a *backup* to the IPv6 address specified with the lms-ipv6 parameter.
LMS Preemption	When this parameter is enabled, the AP automatically reverts to the primary LMS IP address when it becomes available.
LMS Hold-down Period	Time, in seconds, that the primary LMS must be available before an AP returns to that LMS after failover.
GRE Striping IP	Specify an IPv4 address for the .g radio of the controller to allow LACP enabled switches to send traffic for the 2 radios on different links. Recommended value is LMS_IP+1.
Basic AP System Profile Settings–Remote AP	
Remote-AP DHCP Server VLAN	VLAN ID of the remote AP DHCP server used if the controller is unavailable. This VLAN enables the DHCP server on the AP (also known as the remote AP DHCP server VLAN). If you enter the native VLAN ID, the DHCP server is unavailable.

Parameter	Description
Remote-AP DHCP Server Id	IP address used as the DHCP server identifier.
Remote-AP DHCP Default Router	IP address for the default DHCP router.
Remote-AP DHCP DNS Server	IP address of the DNS server.
Remote-AP DHCP Pool Start	Configures a DHCP pool for remote APs. This is the first IP address of the DHCP pool.
Remote-AP DHCP Pool End	Configures a DHCP pool for remote APs. This is the last IP address of the DHCP pool.
Remote-AP DHCP Pool Netmask	Configures a DHCP pool for remote APs. This is the netmask used for the DHCP pool.
Remote-AP DHCP Lease Time	The amount of days that the assigned IP address is valid for the client. Specify the lease in <days>. A value of 0 indicates the IP address is always valid; the lease does not expire.
Remote-AP uplink total bandwidth	This is the total reserved uplink bandwidth (in Kilobits per second).
Remote-AP bw reservation 1 Remote-AP bw reservation 2 Remote-AP bw reservation 3	Session ACLs with uplink bandwidth reservation in kilobits per second. You can specify up to three session ACLs to reserve uplink bandwidth. The sum of the three uplink bandwidths should not exceed the Remote-AP uplink total bandwidth .
Remote-AP Local Network Access	Enable or disable local network access across VLANs in a Remote-AP.
Advanced AP System Profile Settings	
Bootstrap threshold	Number of consecutive missed heartbeats on a GRE tunnel (heartbeats are sent once per second on each tunnel) before an AP rebootstraps. On the controller, the GRE tunnel timeout is 1.5 x bootstrap-threshold; the tunnel is torn down after this number of seconds of inactivity on the tunnel. The supported range is 1-65535, and the default value is 8.
Double Encrypt	This parameter applies only to remote APs. Use double encryption for traffic to and from a wireless client that is connected to a tunneled SSID. When enabled, all traffic is re-encrypted in the IPsec tunnel. When disabled, the wireless frame is only encapsulated inside the IPsec tunnel. All other types of data traffic between the controller and the AP (wired traffic and traffic from a split-tunneled SSID) are always encrypted in the IPsec tunnel.
Dump Server	(For debugging purposes.) Specifies the server to receive a core dump generated when an AP process crashes.
Heartbeat DSCP	Assign a DSCP value to AP heartbeats to prioritize heartbeats traveling over low-speed links. The supported range is 0-63, and the default value is 0. For more information, see Prioritizing AP heartbeats on page 525.
Maintenance Mode	Enable or disable AP maintenance mode. This setting is useful when deploying, maintaining, or upgrading the network.

Parameter	Description
	If enabled, APs stop flooding unnecessary traps and syslog messages to network management systems or network operations centers when deploying, maintaining, or upgrading the network. The controller still generates debug syslog messages if debug logging is enabled.
Number of IPSEC retries	Number of times the AP will try to create an IPsec tunnel with the master controller before the AP will reboot. If you specify a value of 0, and AP will not reboot if it cannot create the IPsec tunnel. The supported range of values is 0-1000 retries, and the default value is 85 retries.
Maximum Request Retries	Maximum number of times to retry AP-generated requests, including keepalive messages. After the maximum number of retries, the AP either tries the IP address specified by the bkup-lms-ip (if configured) or reboots.
Request Retry Interval	Interval, in seconds, between the first and second retries of AP-generated requests. If the configured interval is less than 30 seconds, the interval for subsequent retries is increased up to 30 seconds.
Root AP	Defines a remote AP as the root AP in a branch office network with a multi-AP hierarchy.
AeroScout RTLS Server	Enables the AP to send AeroScout tag information to an RTLS server. You must specify the IP address or DNS server and port number of the server to which location reports are sent. RTLS station reporting includes information for APs and the clients that the AP has detected. If you select the **Include Unassociated Stations** option, the station reports will also include information about clients not associated to any AP. By default, unassociated clients are not included in station reports.
RTLS Server configuration	Enables the AP to send RFID tag information to an RTLS server. You must specify the IP address or DNS server and port number of the server to which location reports are sent, a shared secret key, and the frequency at which packets are sent to the server. RTLS station reporting includes information for APs and the clients that the AP has detected. If you select the **Include Unassociated Stations** option, the station reports will also include information about clients not associated to any AP. By default, unassociated clients are not included in station reports. For more information on configuring RTLS server configuration, see RTLS Server on page 505.
Telnet	Select this checkbox to enable telnet to the AP.
Spanning Tree	Select this checkbox to enable the Spanning Tree protocol.

To configure the bootstrap threshold using the command-line interface, access the CLI in config mode and issue the following command:

```
ap system-profile <profile>
   bootstrap-threshold 30
```

Prioritizing AP heartbeats

If the AP heartbeat or keep alive packets sent between the APs and controller are not received during the defined interval, the APs may reboot, causing clients to re-associate. If a high-latency or low-speed link prevents AP heartbeats from being sent and received correctly, you can assign a DHCP value to AP heartbeats to prioritize the heartbeats.

To prioritize AP heartbeats using the WebUI:

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. Select either the AP Group or AP Specific tab. Click **Edit** for the AP group or AP name.
3. Under Profiles, select **AP**, then **AP system profile**. The configuration settings are displayed in Profile Details.
4. Under Profile Details:
   a. In the **Heartbeat DSCP** field, enter a value greater than zero.
   b. Click **Apply**.

To prioritize AP heartbeats using the command-line interface, access the CLI in config mode and issue the following command:

```
ap system-profile <profile>
   heartbeat-dscp <number>
```

use the following commands:

```
show ap config {ap-group <name>|ap-name <name>|essid <name>}
show ap debug system-status {ap-name <name>|bssid <name>| ip-addr <ipaddr>}
```

On the local controller, you can also view maintenance mode status using the following commands:

```
show ap active {ap-name <name>|essid <name>|ip-addr <ipaddr>}
show ap database
show ap details {ap-name <name>|bssid <name>|ip-addr <ipaddr>}
```

Configuring AP Channel Assignments

The country code in the AP Regulatory Domain profile determines supported channel and channel pairs for that specific AP. Any changes to the country code causes the valid channel lists to be reset to the defaults for the country.

The example in this section illustrates how to perform the following tasks for an AP group:

1. Configure the "default" regulatory domain profile to use a valid country code. (In this example, the country code **US**.) This will determine the available channels.
2. Configure a 40 MHz channel (bonded pair) for the AP group's 802.11a (5 GHz) radio profile.
3. Configure a 20 MHz channel for the AP group's 802.11g (2.4 GHz) radio profile.



This example uses default ARM profile settings and the recommended high-throughput channel assignments for the 802.11a and 802.11b/g bands. If you want the channel assignments to utilize high-throughput, ensure that high-throughput is enabled within the radio profile. For details, see 802.11a and 802.11g RF Management Profiles on page 509.

Using the WebUI

1. Navigate to **Configuration > Wireless > AP Configuration > AP Group** page.
2. Click the **Edit** button by the name of the AP group to which you want to assign specific channels.
3. In the **Profiles** list, expand the AP menu to display the AP profiles used by the AP group.
4. Select the **Regulatory Domain profile** named **default**.
5. Click the **Country Code** drop-down menu and select the **US-United States domain** if it is not already selected.

   The Regulatory Domain's country code determines which channels are selected in the following fields:
   - Valid 802.11g channel
   - Valid 802.11a channel

- Valid 802.11g 40MHz channel pair
- Valid 802.11a 40MHz channel pair

If none of the channels supported by the AP have received regulatory approval by the country whose country code you selected, the AP will revert to Air Monitor mode.

6. In the **Valid 802.11a 80MHz channel group** field, define which 80MHz channels on the 802.11a band are available for assignment by ARM and for the controller to randomly assign if user has not specified a channel. The channel numbers below correspond to channel center frequency.
   - Possible choices in US: 42, 58, 106, 122, 138, 155
   - Possible choices in EU: 42, 58, 106, 122
   - Possible choices in JP: 42, 58, 106, 122
   - Possible global choices: 42, 58, 106, 122, 138, 155
7. Click **Apply**.
8. Under the **Profiles** list, expand the **RF Management** menu.
9. Select the **802.11a radio profile** used by the AP group
10. Enter **36** in the **Channel** text field and select the **Above** radio button. In this instance, channel 36 becomes the primary channel and the secondary channel is 40.
11. Click **Apply**.
12. Under the Profiles list select the **802.11g radio profile** used by the AP group.
13. Enter **1** in the **Channel** text field and select **None**. In this instance, channel 1 is the assigned 20 MHz channel and 40 MHz mode is disabled.
14. Click **Apply**.

Using the CLI

Country codes are generally specified in ISO 3166 format. To see what channels are available for a given country code, use the show ap allowed-channels <country-code> command can be used.

```
ap regulatory-domain-profile default
   country-code US
rf dot11a-radio-profile ht-corpnet-a
   channel 36+
rf dot11g-radio-profile ht-corpnet-g
   channel 1
```



Country codes are generally specified in ISO 3166 format. To see what channels are available for a given country code, use the **show ap allowed-channels country-code <country-code>** command.

Channel Switch Announcement (CSA)

When an AP changes its channel, an existing wireless clients may “time out” while waiting to receive a new beacon from the AP; the client must begin scanning to discover the new channel on which the AP is operating. If the disruption is long enough, the client may need to reassociate, reauthenticate, and request an IP address. Channel Switch Announcement (CSA), as defined by IEEE 802.11h, enables an AP to announce that it is switching to a new channel before it begins transmitting on that channel. This allows the clients, who support CSA, to transition to the new channel with minimal downtime.

When CSA is enabled, the AP does not change to a new channel immediately. Instead, it sends a number of beacons (the default is 4) which contain the CSA announcement before it switches to the new channel. You can configure the number of announcements sent before the change.



Clients must support CSA in order to track the channel change without experiencing disruption.

Using the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. Select either the AP Group or AP Specific tab. Click **Edit** for the AP group or AP name.
3. Select **RF Management** in the Profile list.
4. Select the 802.11a or 802.11g radio profile.
5. Select **Enable CSA**. You can configure a different value for CSA Count.
6. Click **Apply**.

Using the CLI

```
rf radio-profile <profile>
    csa
    csa-count <number>
```

Automatic Channel and Transmit Power Selection

To allow automatic channel and transmit power selection based on the radio environment, enable Adaptive Radio Management (ARM). Note that ARM assignments will override the static channel and power configurations done using the radio profile. For complete information on the Adaptive Radio Management feature, refer to Adaptive Radio Management (ARM) on page 432.

Managing AP Console Settings

An AP's provisioning parameters are unique to each AP. These parameters are initially configured on the controller and then pushed out to the AP and stored on the AP itself. **Best practices are to configure an AP's provisioning settings using the controller WebUI**. If you find it necessary to alter an AP's provisioning settings for troubleshooting purposes, you can do so using the controller WebUI and CLI, or alternatively, through a console connection to the AP itself.

To create a console connection to the AP:

1. Connect a local console to the serial port on the AP. You can connect the AP's serial port to a terminal or terminal server using an Ethernet cable, or connect the serial console port to a DB-9 adapter, then connect the adapter to a laptop using an RS-232 cable. For details on connecting to an AP's serial console port, refer to the Installation Guide included with the AP.
2. Establish a console communication to the AP, then power-cycle the AP to reboot it.
3. To access the AP console command prompt, press **Enter** when the AP displays the message "*Hit <Enter> to stop autoboot.*" If the autoboot countdown expires before you can interrupt it, turn the device off and then back on.
4. Once the AP boot prompt appears, you can issue any of the AP provisioning commands described in the Table 100. Remember, though these commands may be useful for troubleshooting, they are all optional and are *not* necessary for normal AP provisioning.

Table 100: *AP Console Commands*

Command	Description
setenv ipaddr <ipaddr>	IP address to be assigned to the AP.
setenv netmask <netmaskip>	Netmask to be assigned to the AP.
`setenv gatewayip <ipaddr>`	IP address of the internet gateway used by the AP.

Command	Description
`setenv name <ap name>`	Name of the AP.
`setenv group <group name>`	Name of the AP group to which the AP should belong.
setenv master <ipaddr>	IP address of the AP's master controller.
`setenv serverip <ipaddr>`	IP address of the TFTP server from which the AP can download its boot image.
setenv dnsip <ipaddr>	IP address of the DNS server used by the AP.
`setenv domainname <domain>`	Domain name used by the AP.

5. When you are finished, type **Save** and then press **Enter** to save your settings



Other AP console commands may be available when accessing an AP directly through its console port, but these commands can cause configuration errors if used improperly and should only be issued under the direct supervision of Dell technical support.

The example below configures an AP location and domain name using an AP console connection:

```
Hit <Enter> to stop autoboot: 0
apboot> <INTERRUPT>
apboot> setenv group corporate 2
apboot> setenv domainname mycompany.com
apboot> save
apboot>reboot
```

To view current AP settings using the AP console, issue the command **printenv <name>** where **<name>** is one of the variable names listed in Table 100, such as **ipaddr**, **dnsip** or **gatewayip**.

```
apboot> printenv domainname
domainname=mycompany.com
```

Link Aggregation Support on W-AP220 Series and W-AP270 Series

W-AP220 Series and W-AP270 Series access points support link aggregation using either standard port-channel (configuration based) or Link Aggregation Control Protocol (protocol signaling based). These access points can optionally be deployed with LACP configuration to benefit from the higher (greater than 1 Gbps) aggregate throughput capabilities of the two radios.

The controller uses two different IP addresses for different GRE tunnels between the AP and the controller. One IP address is used for tunnels to virtual APs using a 5G radio, while a second controllerIP address is used for tunnels corresponding to virtual APs using a 2.4G radio. The IP addresses should be selected to ensure a different physical interface is used by the load-balancing algorithm on the Ethernet switch. This will allow the W-AP225 and W-AP270 Seriesto achieve greater than 1Gbps throughput in both upstream and downstream directions.

ArubaOS 6.4.2.0 introduces a local AP LACP LMS map information profile that maps a LMS IP address to a GRE striping IP address. If the AP fails over to a standby or backup controller, the AP LACP LMS map information profile on the new controller defines the IP address that AP uses to terminate 802.11g radio tunnels on the new controller. This feature allows W-AP220 Series or W-AP270 Series access points to form a 802.11g radio tunnel to a backup controller the event of a controller failover, even if the backup controller is in a different L3 network.

In previous releases, the GRE striping IP address was defined in the global AP system profile, which did not allow APs to maintain GRE striping tunnels if the AP failed over to a backup controller in a different L3 network.



If your topology includes a backup controller you must define GRE striping IP settings in the active and the backup controller. For more information on LACP features in ArubaOS, see Configuring LACP on page 201.

Configuring LACP

To enable and configure LACP on W-AP220 Series and W-AP270 Series access points, configure the **LMS IP** address and the **GRE Striping IP** address. In ArubaOS 6.4.2 and later, the GRE striping parameter is configured in the AP LACP Striping profile. In ArubaOS 6.3.1.0 to 6.4.1.x, GRE striping is configured in the **AP System profile**. The **GRE Striping IP** value must be an IPv4 address owned by the controller that has the specified **LMS IP**. The **GRE Striping IP** does not belong to any physical or virtual interface on the controller, but the controller can transmit or receive packets using this IP.

You can configure LACP features using the WebUI or the CLI. The procedure varies, depending upon the version of ArubaOS running on your controller.

Using the WebUI, in ArubaOS 6.4.2.x and later

Follow the procedure to configure the LACP parameters in the AP System profile and AP LACP LMS map information profile:

1. Access the active controller and navigate to the **Configuration > Advanced Services > All Profiles** page.
2. Expand the **AP** profiles menu in the **Profiles** pane.
3. Expand the **AP System** profiles menu, and select the AP system profile you want to modify.
4. Select the **Basic** tab on the **Profile Details** pane, locate the **LMS Settings** section and specify a unique IPv4 address in the **LMS IP** field.
5. In the **Profiles** pane, select the **AP LACP LMS map information** profile.
6. In the **Profile Details** pane, select **AP LACP Striping IP** to enable the AP LACP striping feature.
7. Enter a GRE striping IP address in the **IP** field.
8. Enter a LMS IP address in the **LMS** field.
9. Click **Add**.
10. Click **Apply**.

Using the CLI, in ArubaOS 6.4.2.x and later

Execute the following commands to configure LACP and AP LACP LMS map information settings.

```
(host) (config) #ap system-profile LACP
(host) (AP system profile "LACP") #lms-ip 192.0.2.1
(host) (AP system profile "LACP") #exit
(host) (config) #ap-lacp-striping-ip
(host) (AP LACP LMS map information) #striping-ip 192.0.2.2 lms 192.0.2.1
(host) (AP LACP LMS map information) #aplacp-enable
```

Using the WebUI in ArubaOS 6.3.1.x-6.4.1.x

Follow the procedure to configure the LACP parameters in AP System profile:

1. Navigate to the **Configuration > Advanced Services > All Profiles** page.
2. Expand the **AP** profiles menu in the **Profiles** pane
3. Expand the **AP System** profiles menu, and select the AP system profile you want to modify.
4. In the **Profile Details** pane, select the **Basic** tab, locate the **LMS Settings** section and a specify a unique IPv4 address in the **LMS IP** field.

5. Click **Apply**.

Using the CLI in ArubaOS 6.3.1.x-6.4.1.x

Execute the following commands to configure the LACP parameters (LMS IP and the GRE striping IP) on an AP system profile.

```
(host) (config) #ap system-profile LACP
(host) (AP system profile "LACP") #lms-ip 192.0.2.1
(host) (AP system profile "LACP") #gre-striping-ip 192.0.2.2
```

Important Points to Remember

- In the upstream direction when the AP transmits GRE frames to the controller, the bonding driver must be in active-active mode and not in the default active-standby mode to allow link aggregation.
- The W-AP220 Series and W-AP270 Series APs detect the LACP frames and auto-configures itself to LACP mode. If **gre-striping-ip** is not configured, then the AP goes back to the active-standby mode. The AP link may go down in this scenario depending on the behavior of the upstream switch.
- Ensure that the **gre-striping-ip** is unique and not used by any other host on the subnet.
- LACP support is limited to a use case where Enet0 and Enet1 ports of the AP are connected to a switch, and LACP is enabled on the two corresponding switch ports.
- The port priority is not applicable to the AP as both ports need to be used. This value is always set to the maximum numerical priority (0xFF), which is the lowest priority.
- The system priority is not configurable. It is set to the maximum numerical value (0xFFFF), which is the lowest priority. This leaves control of the aggregate to the upstream switch.
- The timeout value is not configurable.
- The key is not configurable and the default key value is 1.
- LACP cannot be enabled if wired AP functionality is enabled on the second port. You cannot enable LACP if the Enet 1 port is shutdown.

Troubleshooting Link Aggregation

The following show commands in the CLI can be used to troubleshoot Link Aggregation on W-AP220 Series and W-AP270 Series APs:

- **show ap debug lacp ap-name <ap-name>**—Using this command, you can view if LACP is active on an AP. It displays the number of GRE packets sent and received on the two Ethernet ports.
- show ap database—TStarting with ArubaOS 6.4.2, the output of this command includes an **LACP Striping** flag (s) to indicate of the AP is configured with a LACP striping IP address,
- **show datapath tunnel**—Using this command, you can verify if the 2.4GHz tunnels are anchored on the **gre-striping-ip** (The GRE IDs for these tunnels are in a range between 0x8300 and 0x83F0)
- **show datapath user**—Using this command, you can verify if the **gre-striping-ip** has an entry with the 'L' (local) flag
- **show datapath route-cache**—Using this command, you can verify if the **gre-striping-ip** has an entry with the controller MAC.

Service Tag

A service tag is a unique seven digit alphanumeric string that is used to electronically identify a Dell device. It is similar to a serial number identifier. Starting with ArubaOS 6.4.2.0, you can view the service tag of some newer

Dell APs from the controller WebUI or CLI. It is displayed along with the serial number in a device information listing.

In the WebUI

To view the service tag of an AP using the WebUI:

1. Navigate to **Configuration > WIRELESS > AP Installation**.
2. View the service tag of an AP in the **AP Service Tag** column of **Provisioning** tab.

In the CLI

Use the following commands to view the service tag of an AP using the CLI:

```
(host) #show ap database long
(host) #show ap details ap-name <ap-name>
(host) #show provisioning-ap-list
```

Use the following commands to rename, regroup, or reprovision an AP with a service tag using the CLI:

```
(host) #ap-rename service-tag <service-tag>
(host) #ap-regroup  service-tag <service-tag>
(host) (config) #provision-ap
(host) (AP provisioning) #reprovision service-tag <service-tag>
```

The secure enterprise mesh solution is an effective way to expand network coverage for outdoor and indoor enterprise environments without any wires. Using mesh, you can bridge multiple Ethernet LANs or you can extend your wireless coverage. As traffic traverses across mesh APs, the mesh network automatically reconfigures around broken or blocked paths. This self-healing feature provides increased reliability and redundancy: the network continues to operate if an AP stops functioning or a connection fails. Dell controllers provide centralized configuration and management for APs in a mesh environment; local mesh APs provide encryption and traffic forwarding for mesh links.

Mesh Overview Information

The following topics in this chapter describes the components of the Dell secure enterprise mesh architecture and profiles, as well as factors that should be taken into consideration when planning your mesh deployment.

- Understanding Mesh Access Points on page 533
- Understanding Mesh Links on page 535
- Understanding Mesh Profiles on page 537
- Understanding Remote Mesh Portals (RMPs) on page 541
- Mesh Deployment Planning on page 545
- Mesh Deployment Solutions on page 543
- Mesh Deployment Planning on page 545

Mesh Configuration Procedures

The following topics describe the procedures required to configure your secure enterprise mesh solution:

1. Creating and Editing Mesh Radio Profiles on page 551
2. Creating and Editing Mesh Radio Profiles on page 551
3. Creating and Editing Mesh High-Throughput SSID Profiles on page 556
4. Configuring Ethernet Ports for Mesh on page 561
5. Provisioning Mesh Nodes on page 564
6. Verifying Your Mesh Network on page 566



Dell strongly recommends staging mesh APs before deploying them. Identify the physical location of the APs, configure them for mesh, provision the APs and verify connectivity before physically deploying them in a live network.

If you are configuring an AP as both a remote access point and a mesh portal, see also Configuring Remote Mesh Portals (RMPs) on page 567

Understanding Mesh Access Points

Mesh APs learn about their environment when they boot up. Mesh APs are either configured as a mesh portal (MPP), an AP that uses its wired interface to reach the controller, or a mesh point (MP), an AP that establishes an all-wireless path to the mesh portal. Mesh APs locate and associate with their nearest neighbor, which

provides the best path to the mesh portal. Mesh portals and mesh points are also known as mesh nodes, a generic term used to describe APs configured for mesh.

A mesh radio's bandwidth can be shared between mesh-backhaul traffic and client traffic. You can, however, configure a radio for mesh services only. If you have a dual-radio AP, a mesh node can be configured to deliver client services on one radio, and both mesh and WLAN services to clients on the other. If you configure a single-radio AP to deliver mesh services only (by disabling the mesh radio in its 802.11a or 802.11g radio profile) that mesh node can not deliver WLAN services to its clients.

For mesh and traditional thin AP deployments, the Dell controller provides centralized provisioning, configuration, policy definition, ongoing network management, and wireless and security services. However, unlike the traditional thin AP case, mesh nodes also perform network traffic encryption and decryption, and packet forwarding over wired and wireless links.

You configure the AP for mesh on the controller using either the WebUI or the CLI. All mesh related configuration parameters are grouped into mesh profiles that you can apply as needed to an AP group or to individual APs.

APs operate as thin APs by default; their primary function is to receive and transmit electromagnetic signals; other WLAN processing is left to the controller. When planning a mesh network, you manually configure APs to operate in mesh portal or mesh point roles. Unlike a traditional WLAN environment, local mesh nodes provide encryption and traffic forwarding for mesh links in a mesh environment. Virtual APs are still applied to non-mesh radios.

Provisioning mesh APs is similar to thin APs; however, there are some key differences. Thin APs establish a channel to the controller from which they receive the configuration for each radio interface. Mesh nodes, in contrast, get their radio interfaces up and running *before* making contact with the controller. This requires a minimum set of parameters from the AP group and mesh cluster so the mesh node discovers a neighbor, and creates a mesh link and subsequent channel with the controller. To do this, you must first define and configure the mesh cluster profile *before* configuring an AP to operate as a mesh node. This chapter first describes how to configure the mesh profile, then describes how to configure APs to operate in mesh mode. If you have already configured a complete mesh profile, continue to "Ethernet Ports for Mesh" or "Provisioning Mesh Nodes".

Mesh Portals

The mesh portal (MPP) is the gateway between the wireless mesh network and the enterprise wired LAN. You configure a Dell AP to perform the mesh portal role, which uses its wired interface to establish a link to the wired LAN. You can deploy multiple mesh portals to support redundant mesh paths (mesh links between neighboring mesh points that establish the best path to the mesh portal) from the wireless mesh network to the wired LAN.

The mesh portal broadcasts the configured mesh service set identifier (MSSID/mesh cluster name), and advertises the mesh network service to available mesh points. Neighboring mesh points that have been provisioned with the same MSSID authenticate to the portal and establish a secure mesh link over which traffic is forwarded. The authentication process requires secure key negotiation, common to all APs, and the mesh link is established and secured using Advanced Encryption Standard (AES) encryption. Mesh portals also propagate channel information, including CSAs.

Mesh Points

The mesh point (MP) is a Dell AP configured for mesh and assigned the mesh point role. Depending on the AP model, configuration parameters, and how it was provisioned, the mesh point can perform multiple tasks. The mesh point provides traditional Dell WLAN services (such as client connectivity, intrusion detection system (IDS) capabilities, user role association, LAN-to-LAN bridging, and Quality of Service (QoS) for LAN-to-mesh communication) to clients and performs mesh backhaul/network connectivity. A mesh radio can be configured

to carry mesh-backhaul traffic only. Additionally, a mesh point can provide LAN-to-LAN Ethernet bridging by sending tagged/untagged VLAN traffic across a mesh backhaul/network to a mesh portal.

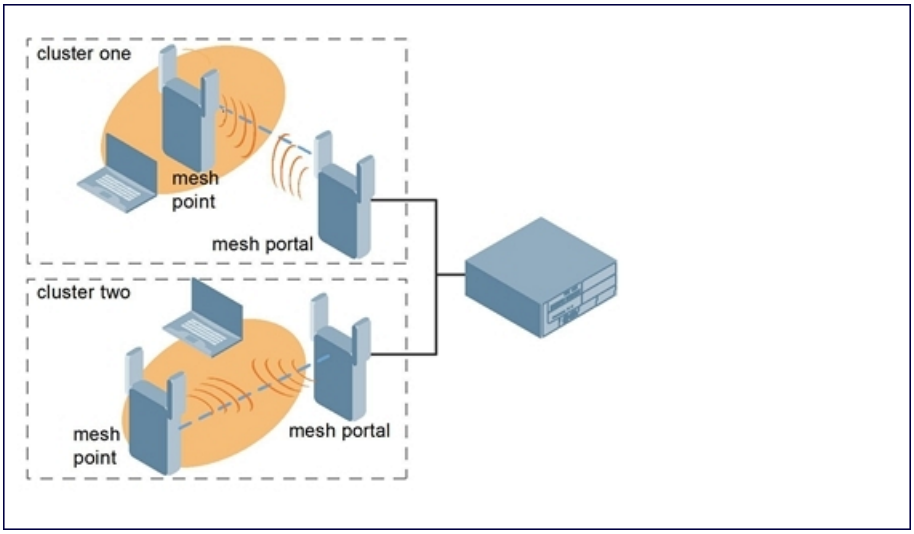
Mesh points use one of their wireless interfaces to carry traffic and reach the controller. Mesh points are also aware of potential neighbors, and can form new mesh links if the current mesh link is no longer preferred or available.

Mesh Clusters

Mesh clusters are similar to an Extended-Service Set (ESS) in a WLAN infrastructure. A mesh cluster is a logical set of mesh nodes that share the common connection and security parameters required to create mesh links. Mesh clusters are grouped and defined by a mesh cluster profile, as described in "Mesh Cluster Profile".

Mesh clusters may enforce predictability in mesh networking by limiting the amount of concurrent mesh points, hop counts, and bandwidth used in the mesh network. A mesh cluster can have multiple mesh portals and mesh points that facilitate wireless communication between wired LANs. Mesh portals in a mesh cluster do not need to be on the same VLAN. Figure 52 shows two mesh clusters and their relationship to the controller.

Figure 52 *Sample Mesh Clusters*



Understanding Mesh Links

The mesh link is the data link between a mesh point and its parent. A mesh point uses the parameters defined in the mesh cluster profile, to establish a mesh link with a neighboring mesh point. The mesh link uses a series of metrics to establish the best path to the mesh portal.

NOTE

Throughout the rest of this chapter, the term "uplink" is used to distinguish the active association between a mesh point and its parent.

The following list describes how mesh links are created:

- Creating the initial mesh link

  When creating the initial mesh link, mesh points look for others advertising the same MSSID as the one contained in its mesh cluster profile. The mesh point scans the channels in its provisioned band of operation to identify a list of neighbors that match its mesh cluster profile. The mesh point then selects the from highest priority neighbors based on the least expected path cost.

  If no provisioned mesh cluster profile is available, mesh points use the recovery profile to establish an uplink. If multiple cluster profiles are configured, mesh points search, in order of priority, their list of

provisioned backup mesh cluster profiles to establish an uplink. If the configured profiles are unavailable after searching for 5 minutes, the recovery profile is used.

- Moving to a better mesh link

  If the existing uplink quality degrades below the configured threshold, and a lower cost or more preferable uplink is available on the same channel and cluster, the mesh point reselects that link without re-scanning. In some cases, this invalidates all of the entries that have this mesh point as a next hop to the destination and triggers new learning of the bridge tables.

- Using a new mesh link if the current mesh link goes down

  If an uplink goes down, the affected mesh nodes reestablish a connection with the mesh portal by re-scanning to choose a new path to the mesh portal. If a mesh portal goes down, and a redundant mesh portal is available, the affected mesh nodes update their forwarding tables to reflect the path to the new mesh portal.

Link Metrics

Mesh points use the configured algorithm to compute a metric value, or "path cost," for each potential uplink and select the one with the lowest value as the optimal path to the mesh portal. Table 101 describes the components that make up the metric value: node cost, hop count, link cost and 802.11 capacity.

The link metrics indicate the relative cost of a path to the mesh portal. The best path (lowest metric value) is used to create the uplink.

Table 101: *Mesh Link Metric Computation*

Component	Description
Node cost	Indicates the amount of traffic expected to traverse the mesh node. The more traffic, the higher the node cost. When establishing a mesh link, nodes with less traffic take precedence. The node cost is dependent on the number of children a mesh node supports. It can change as the mesh network topology changes, for example if new children are added to the network or old children disconnect from the network.
Hop count	Indicates the number of hops it takes the mesh node to get to the mesh portal. The mesh portal advertises a hop count of 0, while all other mesh nodes advertise a cumulative count based on the parent mesh node.
Link cost	Represents the quality of the link to an active neighbor. The higher the Received Signal Strength Indication (RSSI), the better the path to the neighbor and the mesh portal. If the RSSI value is below the configured threshold, the link cost is penalized to filter marginal links. A less direct, higher quality link may be preferred over the marginal link. The following factors also affect mesh link metrics: • High-throughput APs add a high cost penalty for links to non-high-throughput APs. • Multi-stream high-through APs add proportional cost penalties for links to high-throughput APs that support fewer streams.
802.11 capacity	High-throughput APs can send 802.11 information elements (IEs) in their management frames, allowing high-throughput mesh nodes to identify other mesh nodes with a high-throughput capacity. High-throughput mesh points prefer to select other 802.11-capable mesh points in their path to the mesh portal, but can use a legacy path if no high-throughput path is available.
Path Cost	Path cost is calculated by analyzing the other components in this table, and adding the link cost, the mesh parent's path cost, and the parent's node cost. Mesh portals typically advertise a path-cost of zero, but high-throughput portals add an offset penalty if they are connected to a 10/100mbps port that is too slow for the high-throughput link capacity.

Optimizing Links

You can configure and optimize operation of the link metric algorithm via the mesh radio profile. These configurable mesh link trigger thresholds can determine when the uplink or mesh path is dropped and another is chosen, provide enhanced network reliability, and contain flapping links. Although you can modify the behavior of the link metric algorithm, It is recommended to follow the default values for most deployments. For information, see Metric algorithm on page 553.

Understanding Mesh Profiles

Mesh profiles help define and bring-up the mesh network. The following sections describe the mesh cluster, mesh radio, and mesh recovery profiles in more detail.

The complete mesh profile consists of a mesh radio profile, RF management (802.11a and 802.11g) radio profiles, a high-throughput SSID profile (if your deployment includes 802.11n-capable APs), a mesh cluster profile, and a read-only recovery profile. The recovery profile is dynamically generated by the master controller; you do not explicitly configure the recovery profile.

Dell provides a "default" version of the mesh radio, RF management, high-throughput SSID and cluster profiles with default values for most parameters. You can use the "default" version of a profile or create a new instance of a profile which you can then edit as you need. You can change the values of any parameter in a profile. You have the flexibility of applying the "default" versions of profiles in addition to customizing profiles that are necessary for the AP or AP group to function.

If you assign a profile to an individual AP, the values in the profile override the profile assigned to the AP group to which the AP belongs. The exception is the mesh cluster profile: you can apply multiple mesh cluster profiles to individual APs, as well as to AP groups.

Mesh Cluster Profiles

Mesh clusters are grouped and defined by a mesh cluster profile, which provides the framework of the mesh network. Similar to virtual AP profiles, the mesh cluster profile contains the MSSID (mesh cluster name), authentication methods, security credentials, and cluster priority required for mesh nodes to associate with their neighbors and join the cluster. Associated mesh nodes store this information in flash memory. Although most mesh deployments require only a single mesh cluster profile, you can configure and apply multiple mesh cluster profiles to an AP group or an individual AP. If you have multiple cluster profiles, the mesh portal uses the profile with the highest priority to bring up the mesh network. Mesh points, in contrast, go through the list of mesh cluster profiles in order of priority to decide which profile to use to associate themselves with the network. The mesh cluster priority determines the order by which the mesh cluster profiles are used. This allows you, rather than the link metric algorithm, to explicitly segment the network by defining multiple cluster profiles.

Since the mesh cluster profile provides the framework of the mesh network, you must define and configure the mesh cluster profile before configuring an AP to operate as a mesh node. You can use either the **default** cluster profile or create your own. If you find it necessary to define more than one mesh cluster profile, you must assign priorities to each profile to allow the Mesh AP group to identify the primary and backup mesh cluster profile(s). The primary mesh cluster profile and each backup mesh cluster profile must be configured to use the same RF channel. The APs may not provision correctly if they are assigned to a backup mesh cluster profile with a different RF channel than the primary mesh cluster profile.

If the mesh cluster profile is unavailable, the mesh node can revert to the recovery profile to bring-up the mesh network until the cluster profile is available. You can also exclude one or more mesh cluster profiles from an individual AP—this prevents a mesh cluster profile defined at the AP group level from being applied to a specific AP.

Do not delete or modify mesh cluster profiles once you use them to provision mesh nodes. You can recover the mesh point if the original cluster profile is still available. It is recommended to create a new mesh cluster profile if needed. If you modify any mesh cluster setting, you must reprovision your AP for the changes to take effect (this also causes the AP to automatically reboot). See "Provisioning Mesh Nodes" for more information.

If you configure multiple cluster profiles with different cluster priorities, you manually override the link metric algorithm because the priority takes precedence over the path cost. In this scenario, the mesh portal uses the profile with the highest priority to bring-up the mesh network. The mesh portal stores and advertises that one profile to neighboring mesh nodes to build the mesh network. This profile is known as the "primary" cluster profile. Mesh points, in contrast, go through the list of configured mesh cluster profiles in order of priority to find the profile being advertised by the mesh portal. Once the primary profile has been identified, the other profiles are considered "backup" cluster profiles. Use this deployment if you want to enforce a particular mesh topology rather than allowing the link metric algorithm to determine the topology.

For this scenario, do the following:

- Configure multiple mesh cluster profiles with different priorities. The primary cluster profile has a lower priority number, which gives it a higher priority.
- Configure the mesh radio profile.
- Create an AP group for 802.11a radios and 802.11g radios
- Configure the 802.11a or 802.11g RF management profiles for each AP group.
- If your deployment includes high-throughput APs, configure the mesh high-throughput SSID profile. The mesh radio profile uses the default high-throughput SSID profile unless you specifically configure the mesh radio profile to use a different high-throughput SSID profile
- Create an AP group for each 802.11a channel.

If a mesh link breaks or the primary cluster profile is unavailable, mesh nodes use the highest priority backup cluster profile to re-establish the uplink or check for parents in the backup profiles. If these profiles are unavailable, the mesh node can revert to the recovery profile to bring up the mesh network until a cluster profile is available. For information about the procedure to configure a mesh cluster profile, see Configuring Mesh Cluster Profiles on page 547

Mesh Radio Profiles

The mesh radio profile allows you to specify the set of rates used to transmit data on the mesh link. This profile also allows you to define a **reselection-mode** setting to optimize the operation of the link metric algorithm. The reselection mode specifies the method a mesh node uses to find a better uplink to create a path to the mesh portal. Only neighbors on the same channel in the same mesh cluster are considered.

The mesh radio profile includes the following reselection mode options:

- **reselect-anytime**: mesh points using the **reselect-anytime** reselection mode perform a single topology readjustment scan within 9 minutes of startup and 4 minutes after a link is formed. If no better parent is found, the mesh point returns to its original parent. This initial scan evaluates more distant mesh points before closer mesh points, and incurs a dropout of 5–8 seconds for each mesh point. After the initial startup scan is completed, connected mesh nodes evaluate mesh links every 30 seconds. If a mesh node finds a better uplink, the mesh node connects to the new parent to create an improved path to the mesh portal.
- **reselect-never**: connected mesh nodes do not evaluate other mesh links to create an improved path to the mesh portal.
- **startup-subthreshold**: mesh points using the **startup-subthreshold** reselection mode perform a single topology readjustment scan within 9 minutes of startup and 4 minutes after a link is formed. If no better parent is found, the mesh point returns to its original parent. This initial startup scan evaluates more distant mesh points before closer mesh points, and incurs a dropout of 5–8 seconds for each mesh point. After

that time, each mesh node evaluates alternative links if the existing uplink falls below the configured threshold level (the link becomes a sub-threshold link). It is recommended to use this default **startup-subthreshold** value.

- **subthreshold-only**: connected mesh nodes evaluate alternative links only if the existing uplink becomes a sub-threshold link.

If a mesh point using the **startup-subthreshold** or **subthreshold-only** mode reselects a more distant parent because its original, closer parent falls below the acceptable threshold, then as long as that mesh point is connected to that more distant parent, it seeks to reselect a parent at the earlier, shorter distance (or less) with good link quality. For example, if a mesh point disconnects from a mesh parent 2 hops away and subsequently reconnects to a mesh parent 3 hops away, then the mesh point continues to seek a connection to a mesh parent with both an acceptable link quality and a distance of two hops or less, even if the more distant parent also has an acceptable link quality.

For information about the procecure to configure mesh radio profiles, see Creating and Editing Mesh Radio Profiles on page 551.

RF Management (802.11a and 802.11g) Profiles

The two 802.11a and 802.11g RF management profiles for an AP configure its 802.11a (5 Ghz) and 802.11b/g (2.4 GHz) radio settings. Use these profile settings to determine the channel, beacon period, transmit power, and ARM profile for a mesh AP's 5 GHz and 2.5 Ghz frequency bands. You can either use the "default" version of each profile, or create a new 802.11a or 802.11g profile which you can then configure as necessary. Each RF management profile also has a **radio-enable** parameter that allows you to enable or disable the AP's ability to simultaneously carry WLAN client traffic and mesh-backhaul traffic on that radio. This value is enabled by default. For information about configuring RF Management Radio profiles, see 802.11a and 802.11g RF Management Profiles on page 509.



If you do not want the mesh radios carrying mesh-backhaul traffic to support client traffic, consider using a dedicated 802.11a/80211/g radio profile with the mesh radio disabled. In this scenario, the radio carries mesh backhaul traffic but does not support client Virtual APs.

Mesh nodes operating in different cluster profiles can share the same radio profile. Conversely, mesh portals using the same cluster profile can be assigned different RF Management Radio profiles to achieve frequency separation (for more information, see "Deployments with Multiple Mesh Cluster Profiles").

Adaptive Radio Management Profiles

Each 802.11a and 802.11g radio profile references an Adaptive Radio Management (ARM) profile. When you assign an active ARM profile to a mesh radio, ARM's automatic power-assignment and channel-assignment features automatically select the radio channel with the least amount of interference for each mesh portal, maximizing end user performance. In earlier versions of this software, an AP with a mesh radio received its beacon period, transmission power, and 11a/11g portal channel settings from its mesh radio profile. Mesh-access AP portals now inherit these radio settings from their dot11a or dot11g radio profiles.

Each ARM-enabled mesh portal monitors defined thresholds for interference, noise, errors, rogue APs and radar settings, then calculates interference and coverage values and selects the best channel for its radio band (s). The mesh portal communicates its channel selection to its mesh points via Channel Switch Announcements (CSAs), and the mesh points change their channel to match their mesh portal. Although channel settings can still be defined for a mesh point via that mesh point's 802.11a and 802.11g radio profiles, these settings are overridden by any channel changes from the mesh portal. A mesh point takes the same channel setting as its mesh portal, regardless of its associated clients. If you want to manually assign channels to mesh portals or

mesh points, disable the ARM profile associated with the 802.11a or 802.11g radio profile by setting the ARM profile's **assignment** parameter to **disable**.

Mesh points, unlike mesh portals, do not scan channels. This means that once a mesh point has selected a mesh portal or an upstream mesh point, it tunes to this channel, forms the link, and does not scan again unless the mesh link gets broken. This provides good mesh link stability, but may adversely affect system throughput in networks with mesh portals and mesh points. When ARM assigns optimal channels to mesh portals, those portals use different channels. Once the mesh network has formed and all the mesh points have selected a portal (or upstream mesh point), those mesh points are not be able to detect other portals on other channels that could offer better throughput. This type of suboptimal mesh network may form if, for example, two or three mesh points select the same mesh portal after booting, form the mesh network, and leave a nearby mesh portal without any mesh points. Again, this does not affect mesh functionality, but may affect total system throughput. For details about associating an ARM profile with a mesh AP, see Assigning an ARM Profile on page 516.

High-Throughput Radio Profiles

Each 802.11a and 802.11g radio profile also references a high-throughput profile that manages an AP or AP group's 40Mhz tolerance settings. For information about referencing a high-throughput profile, see Assigning a High-throughput Profile on page 515.

Mesh High-Throughput SSID Profiles

High-throughput APs support additional settings not available in legacy APs. A mesh high-throughput SSID profile can enable or disable high-throughput (802.11n) features and 40 MHz channel usage, and define values for aggregated MAC protocol data units (MDPUs) and Modulation and Coding Scheme (MCS) ranges.

Dell provides a "default" version of the mesh high-throughput SSID profile. You can use the "default" version or create a new instance of a profile which you can then edit as you need. High-throughput mesh nodes operating in different cluster profiles can share the same high-throughput SSID radio profile. For information about configuring mesh high-throughput SSID profiles, see "Mesh High-Throughput SSID Profiles".

Wired AP Profiles

The wired AP profile controls the configuration of the Ethernet port(s) on your AP. You can use the wired AP profile to configure Ethernet ports for bridging or secure jack operation using the wired AP profile. For details, see Configuring Ethernet Ports for Mesh on page 561

Mesh Recovery Profiles

In addition to the "default" and user-defined mesh cluster profiles, mesh nodes also have a recovery profile. The master controller dynamically generates a recovery profile, and each mesh node provisioned by the same master controller has the same recovery profile. The recovery profile is based on a pre-shared key (PSK), and mesh nodes use the recovery profile to establish a link to the controller if the mesh link is broken and no other mesh cluster profiles are available.

The mesh portal advertises the provisioned cluster profile. If a mesh point is unaware of the active mesh cluster profile, but is aware of and has the same recovery profile as the mesh portal, the mesh point can use the recovery profile to connect to the mesh portal.

The mesh point must have the same recovery profile as the parent to which it connects. If you provision the mesh points with the same master controller, the recovery profiles should match.



To verify that the recovery profile names match, use the following command:
`show ap mesh debug provisioned-clusters {ap-name <name> | bssid <bssid> | ip-addr <ipaddr>}.`

To view the recovery profile on the controller, use the following command:
`show running-config | include recovery.`

If a mesh point connects to a parent using the recovery profile, it may immediately exit recovery if the parent is actively using one of its provisioned mesh cluster profiles. Once in recovery, a mesh point periodically exits recovery to see if it can connect using an available provisioned mesh cluster profile. The recovery profile is read-only; it cannot be modified or deleted.

The recovery profile is stored in the master controller's configuration file and is unique to that master controller. If necessary, you can transfer your configuration to another controller. If you do so, make sure your new mesh cluster is running and you have re-provisioned the mesh nodes before deleting your previous configuration. The APs learn the new recovery profile after they are provisioned with the new controller. This is also true if you provision a mesh node with one master controller and use it with a different master controller. In this case, the recovery profile does not work on the mesh node until you re-provision it with the new master controller.

Understanding Remote Mesh Portals (RMPs)

You can deploy mesh portals to create a hybrid mesh/remote AP environment to extend network coverage to remote locations; this feature is called remote mesh portal, or RMP. The RMP feature integrates the functions of a remote AP (RAP) and the Mesh portal. As a RAP, it sets up a VPN tunnel back to the corporate switch that secures control traffic between the RAP and the switch.

The Remote Mesh Portal feature allows you to configure a remote AP at a branch office to operate as a mesh portal for a mesh cluster. Other mesh points belonging to that cluster get their IP address and configuration settings from the main office via an IPsec tunnel between the remote mesh portal and the main office controller. This feature is useful for deploying an all-wireless branch office or creating a complete wireless network in locations where there is no wired infrastructure in place.

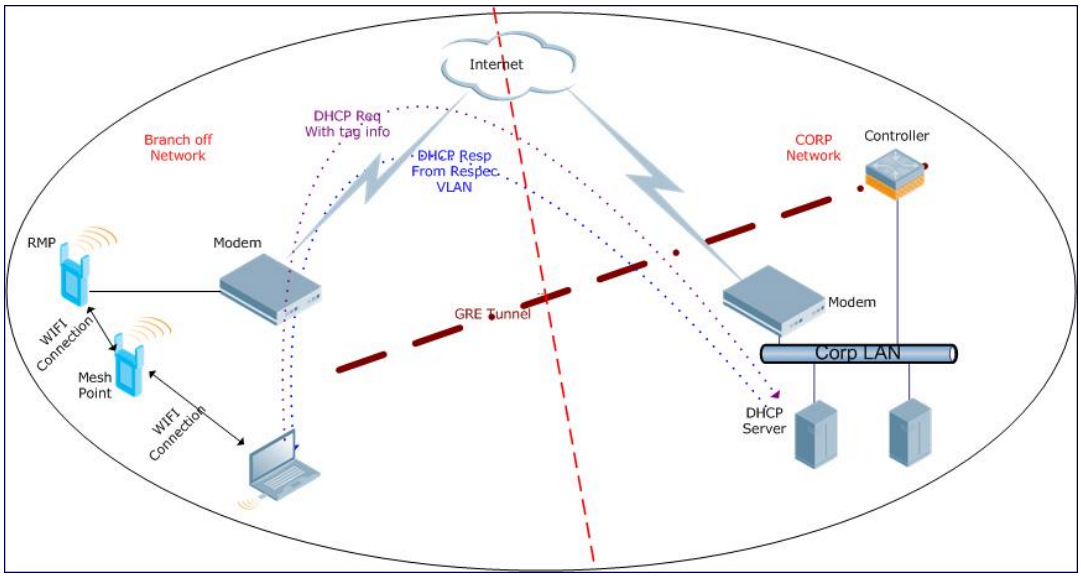
When the client at the branch office associates to a virtual AP in split-tunnel forwarding mode, the client's DHCP requests are forwarded over a GRE tunnel (split tunnel) to the corporate network. This communication is done over a secure VPN tunnel. The IPs are assigned from the corporate pool based on the VLAN tag information, which helps to determine the corresponding VLAN. The VLAN tag also determines the subnet from which the DHCP address has assigned.

A mesh point sends the DHCP request with the mesh private VLAN (MPV) parameter. The mesh point learns the MPV value from the response during the mesh association. When the split tunnel is setup for the RMP on the controller, the VLAN of the tunnel should be the MPV.A DHCP pool for the MPV should be setup on the switch. The use of MPV makes it easy for the RMP to decide which requests to forward over the split tunnel. All requests tagged with the MPV are sent over the split tunnel. Hence the MPV should be different from any user VLAN that is bridged using the mesh network.



The RMP configuration requires an AP license. For more information about Dell software licenses, see Software Licenses on page 130.

Figure 53 *Working of RMP*



By default, the data frames the mesh portal receives on its mesh link are forwarded according to the bridge table entries on the portal. However, frames received on mesh private VLAN (MPV) are treated differently by the remote mesh portal. These frames are treated the same as frames received on a split SSID and are routed rather than bridged. Mesh points obtain DHCP addresses from the corporate network. then register with the controller using these IP addresses. When these mesh points send and receive PAPI control traffic from the main office controller, it controls these mesh points just as if they were on a local VLAN. PAPI traffic containing keys and other secret information receives IPsec encryption and decryption when it is forwarded to the controller through the VPN tunnel.

Not all traffic from a mesh point is sent on the mesh private VLAN. When a mesh point bridges data received via its Ethernet interface or from clients connected to an access radio VAP, the mesh point does not tag the frame with the mesh private VLAN tag when it sends the data through mesh link to the remote mesh portal. Note that the mesh point may still tag the frame depending on the VLAN of the virtual AP and the native VLAN specified in the system profile. Care must be taken to assign the MPV value so that it does not clash with any local tags assigned in the mesh network. In this scenario, the portal performs the default operation and bridges the frame based on its bridge table. Traffic destined to the Internet is recognized as such by the remote mesh portal based on ACL rules. This traffic is NATed on the remote mesh portal's Ethernet interface.

For information on the procedure to configure remote mesh portals, see Configuring Remote Mesh Portals (RMPs) on page 567

Understanding the AP Boot Sequence

The section describes the boot sequence for mesh APs in detail. Depending on its configured role, the AP performs a slightly different boot sequence.

Booting the Mesh Portal

When the mesh portal boots, it recognizes that one radio is configured to operate as a mesh portal. It then obtains an IP address from a DHCP server on its Ethernet interface, discovers the master controller on that interface, registers the mesh radio with the controller, and obtains regulatory domain and mesh radio profiles for each mesh point interface. A mesh virtual AP is created on the mesh portal radio interface, the regulatory

domain and radio profiles are used to bring up the radio on the correct channel, and the provisioned mesh cluster profile is used to setup the mesh virtual AP with the correct announcements on beacons and probe responses. On the non-mesh radio provisioned for access mode, that radio is a thin AP and everything on that interface works as a thin AP radio interface.

If the 802.11a/802.11g radio profile assigned to the mesh radio is enabled, the radio supporst both mesh backhaul and client access Virtual APs. If the mesh radio is to be used exclusively for mesh backhaul traffic, associate that radio to a dedicated 802.11a/802.11g radio profile with the radio disabled so the mesh radios carry backhaul traffic only.

Booting the Mesh Point

When the mesh point boots, it scans for neighboring mesh nodes to establish a link to the mesh portal. All of the mesh nodes that establish the link are in the same mesh cluster. After the link is up, the mesh point uses the DHCP to obtain an IP address and uses the same master controller as their parent. The remaining boot sequence, if applicable, is similar to that of a thin AP. Remember, the priority of the mesh point is establishing a link with neighboring mesh nodes, not establishing a control link to the controller.



In a single hop environment, the mesh point establishes a direct link with the mesh portal.

Air Monitoring and Mesh

Each mesh node has an air monitor (AM) process that registers the BSSID and the MAC address of the mesh node to distinguish it from a thin AP. This allows the WLAN management system (WMS) on the controller and AMs deployed in your network to distinguish between APs, wireless clients, and mesh nodes. The WMS tables also identify the mesh nodes.

For all thin APs and mesh nodes, the AM identifies a mesh node from other packets monitored on the air, and the AM does not trigger wireless-bridging events for packets transmitted between mesh nodes.

Mesh Deployment Solutions

You can configure the following single-hop and multi-hop solutions:

- Thin AP services with wireless backhaul deployment
- Point-to-point deployment
- Point-to-multipoint deployment
- High-availability deployment

With a thin AP wireless backhaul deployment, mesh provides services and security to remote wireless clients and sends all control and user traffic to the master controller over a wireless backhaul mesh link.

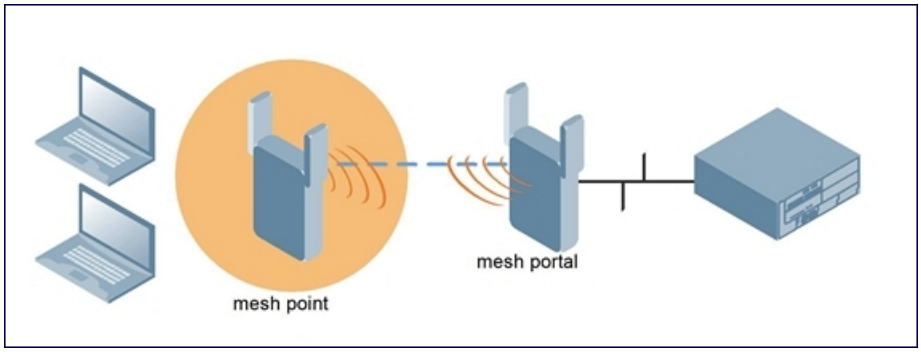
The remaining deployments allow you to extend your existing wired network by providing a wireless bridge to connect Ethernet LAN segments. You can use these deployments to bridge Ethernet LANs between floors, office buildings, campuses, factories, warehouses, and other environments where you do not have access to physical ports or cable to extend the wired network. In these scenarios, a wireless backhaul carries traffic between the Dell APs configured as the mesh portal and the mesh point, to the Ethernet LAN.

Thin AP Services with Wireless Backhaul Deployment

To expand your wireless coverage without bridging Ethernet LAN segments, you can use thin AP services with a wireless backhaul. In this scenario, the mesh point provides network access for wireless clients and establishes a mesh path to the mesh portal, which uses its wired interface to connect to the controller. Use the 802.11g

radio for WLAN and controller services and the 802.11a radio for mesh services. Figure 54 shows the wireless backhaul between the mesh portal to the mesh point that services the wireless clients.

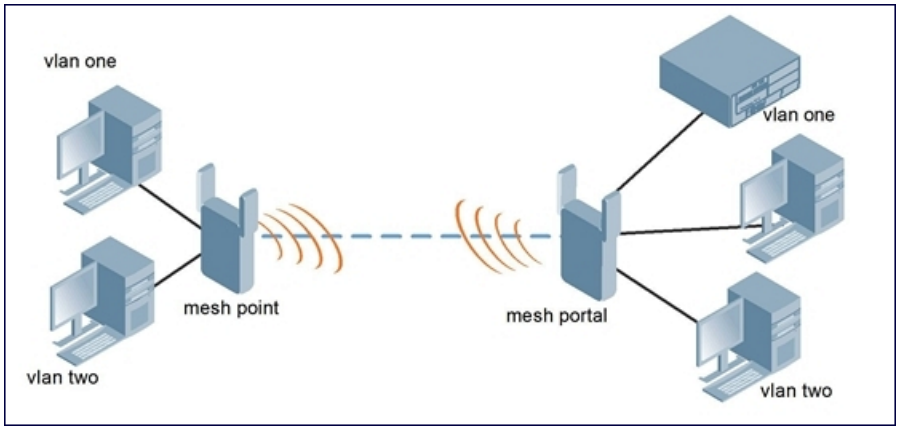
Figure 54 *Sample Wireless Backhaul Deployment*



Point-to-Point Deployment

In this point-to-point scenario, two Ethernet LAN segments are bridged via a wireless connection that carries both client services traffic and mesh-backhaul traffic between the mesh portal and the mesh point. This provides communication from one LAN to another. Figure 55 shows a single-hop point-to-point deployment.

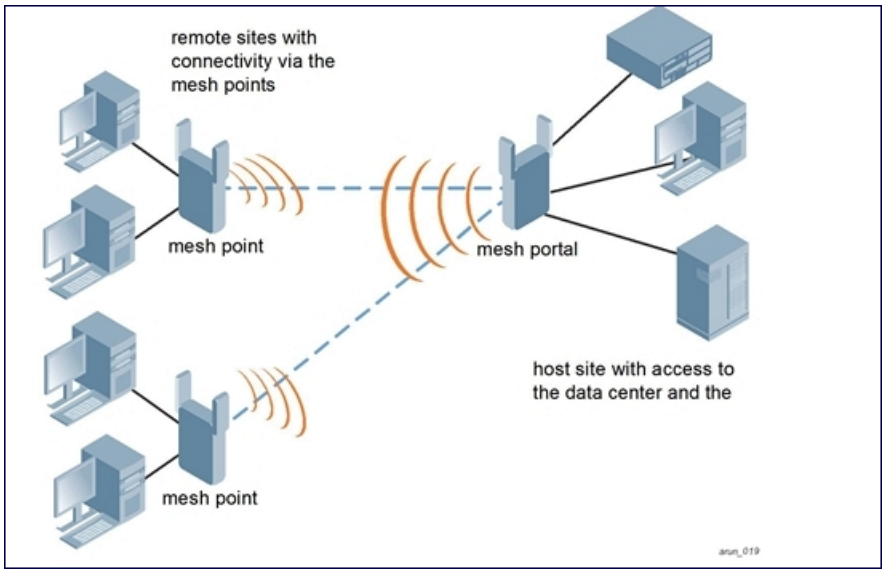
Figure 55 *Sample Point-to-Point Deployment*



Point-to-Multipoint Deployment

In a point-to-multipoint scenario, multiple Ethernet LAN segments are bridged via multiple wireless/mesh backhauls that carry traffic between the mesh portal and the mesh points. This provides communication from the local LAN to multiple remote LANs. Figure 56 shows a single-hop point-to-multipoint deployment.

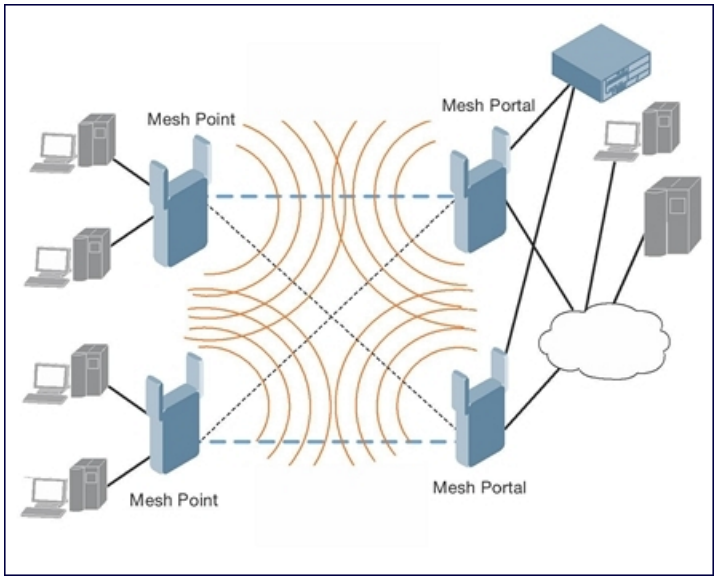
Figure 56 *Sample Point-to-Multipoint Deployment*



High-Availability Deployment

In this high-availability scenario, multiple Ethernet LAN segments are bridged via multiple wireless backhauls that carry traffic between the mesh portal and the mesh points. You configure one mesh portal for each remote LAN that you are bridging with the host LAN. This provides communication from the host LAN to multiple remote LANs. In the event of a link failure between a mesh point and its mesh portal, the affected mesh point could create a link to the other mesh portal. Figure 57 shows a sample single-hop high-availability deployment. The dashed lines represent the current mesh link between the mesh points and their mesh portals. The diagonal dotted lines represent possible links that could be formed in the event of a mesh link or mesh portal failure.

Figure 57 *Sample High-Availability Deployment*



Mesh Deployment Planning

Following considerations are recommended when planning and deploying a mesh solution:

Pre-Deployment Considerations

- Stage the APs before deployment. Identify the location of the APs, configure them for mesh, provision them, and verify connectivity before physically deploying the mesh APs in a live network.
- Ensure the controller has Layer-2/3 network connectivity to the network segment where you plan to install the mesh portal.
- Keep the AP packaging materials and reuse them to send the APs to the installation location.
- Verify the layout of the physical location to determine the appropriate configuration and placement of the APs. Use this information to avoid problems that would necessitate a physical recovery.
- Label the AP before sending it to the physical location for installation.

Outdoor-Specific Deployment Considerations

- Provision the AP with the latitude and longitude coordinates of the installation location. This allows you to more easily identify the AP for inventory and troubleshooting purposes.
- Identify a "radio line of sight" between the antennas for optimum performance. The radio line of sight involves the area along a link through which the bulk of the radio signal power travels.
- Identify the minimum antenna height required to ensure a reliable mesh link.
- Scan your proposed site to avoid radio interference caused by other radio transmissions using the same or an adjacent frequency.
- Consider extreme weather conditions known to affect your location, including: temperature, wind velocity, lightning, rain, snow, and ice.
- Allow for seasonal variations, such as growth of foliage.

For more detailed outdoor deployment information, refer to the installation guide that came with your outdoor AP.

Configuration Considerations

- On dual-radio APs, you can configure only one of the radio for mesh. If you want a dual-radio AP to carry mesh backhaul traffic and client services traffic on separate radios, It is recommended to use 802.11a radios for mesh-backhaul traffic and 802.11g radios for traditional WLAN access.
- If you configure more than one mesh node in the same VLAN, prevent network loops by enabling STP on the Layer-2 switch used to connect the mesh nodes.
- Mesh nodes learn a maximum of 1,024 source MAC addresses; this cannot be changed.
- Place all APs for a specific mesh cluster in the same AP group.
- Create and keep separate mesh cluster profiles for specific mesh clusters. Do not overwrite or delete the cluster profiles.
- Enable bridging on mesh point Ethernet ports when deploying LAN bridging solutions.
- APs configured as mesh points support secure jack operation on enet0. APs with multiple Ethernet ports configured as mesh *portals* support secure jack operation on enet1. If an AP with multiple Ethernet ports is configured as a mesh *point*, it supports secure jack operation on enet1 and enet0.
- Mesh networks forward tagged/untagged VLAN traffic, but do not tag traffic. The allowed VLANS are controlled by the wired ap profile.
- Mesh APs provisioned on different controllers can interoperate if those APs are configured with the same country code, cluster name and cluster key.However, the mesh recovery profile created on one controller is not able to recover settings for mesh APs provisioned on another controller unless the recovery profile is on a master controller and the other mesh nodes were provisioned by a local controller connected to that master.

Post-Deployment Considerations

- Do not connect mesh point Ethernet ports in such a way that causes a network loop.
- Have a trained professional install the AP. After installation, check to ensure the AP receives power and boots up, enabling RSSI outputs.



Although the AP is up and operational, it is not connected to the network.

- Align the AP antenna for optimal RSSI.
- Do not delete or modify mesh cluster profiles once you use them to provision mesh nodes. You can recover the mesh point if the original cluster profile is still available. It is recommended to create a new mesh cluster profile if needed.
- If you create a new mesh cluster profile for an existing deployment, you must re-provision the AP for the new profile to take effect. If you re-provision mesh nodes that are already operating, re-provision the most distant (highest hop count) mesh points first, followed by the mesh portals. If you re-provision the mesh portal first, the mesh points may be unable to form a mesh link. Note that re-provisioning the AP causes it to automatically reboot, which may cause a disruption of service to the network.

Dual-Port AP Considerations

A dual-port AP has two 10/100 Mbps Ethernet ports (enet0 and enet1, respectively). When using these APs in a mesh environment, note the following Ethernet port requirements:

- If configured as a mesh portal:
  - Connect enet0 to the controller to obtain an IP address. The wired AP profile controls enet1.
  - Only enet1 supports secure jack operation.
- If configured as a mesh point, enet0 and enet1 can be configured using separate wired-port-profiles

Configuring Mesh Cluster Profiles

The mesh cluster configuration gets pushed from the controller to the mesh portal and the other mesh points, which allows them to inherit the characteristics of the mesh cluster of which they are a member. Mesh nodes are grouped according to a mesh cluster profile that contains the MSSID, authentication methods, security credentials, and cluster priority. Cluster profiles (including the **default** cluster profile) are not applied until you provision your APs for mesh. For more information on mesh cluster profiles, see Mesh Cluster Profiles on page 537

Managing Mesh Cluster Profiles in the WebUI

Use the following procedures to define and manage mesh cluster profiles using the WebUI.

Creating a Profile

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select the **AP Group** or **AP Specific** tab.
   - If you selected **AP Group**, click **Edit** by the AP group name for which you want to create the new mesh cluster profile.
   - If you selected **AP Specific**, click **Edit** by AP for which you want to create the new mesh cluster profile.
2. In the **Profiles** list, expand the **Mesh** menu, then select **Mesh Cluster profile**.
3. In the **Profile Details** window pane, click the **Add a profile** drop-down list and select **NEW**.
4. Enter a name for the new profile.

5. Configure the mesh cluster settings described in Table 102, then click **Apply**.

Table 102: *Mesh Cluster Profile Configuration Parameters*

Parameter	Description
Profile Name	Name of the mesh cluster profile. The name must be 1–63 characters. Default: Mesh cluster profile named "default."
Cluster Name	Indicates the mesh cluster name. The name can have a maximum of 32 characters, and is used as the MSSID for the mesh cluster. When you first create a new mesh cluster profile, the profile uses the default cluster name "Dell-mesh". Use the Cluster Name parameter to define a new, unique MSSID before you assign APs or AP groups to the mesh cluster profile. **NOTE:** If you want a mesh cluster to use WPA2-PSK-AES encryption, do not use spaces in the mesh cluster name, as this may cause errors in mesh points associated with that mesh cluster. To view existing mesh cluster profiles, use the CLI command: show ap mesh-cluster-profile. A mesh portal chooses the best cluster profile and provisions it for use. A mesh point can have a maximum of 16 cluster profiles. Default: Mesh cluster named "Dell-mesh."
RF Band	Indicates the band for mesh operation for multiband radios. Select a or g. Important: If you create more than one mesh cluster profile for an AP or AP group, each mesh cluster profile must use the same band.
Encryption	Configures the data encryption, which can be either opensystem (no authentication or encryption) or wpa2-psk-aes (WPA2 with AES encryption using a preshared key). It is recommended to select wpa2-psk-aes and using the wpa-passphrase parameter to select a passphrase. Keep the passphrase in a safe place. Default: opensystem.
WPA Hexkey	Configures a WPA pre-shared key. This key must be 64 hexadecimal characters
WPA Passphrase	Sets the WPA password that generates the PSK. The passphrase must be between 8–63 characters, inclusive.
Priority	Indicates the priority of the cluster profile. The mesh cluster priority determines the order by which the mesh cluster profiles are used. This allows you, rather than the link metric algorithm, to control the network topology by defining the cluster profiles to use if one becomes unavailable Specify the cluster priority when creating a new profile or adding an existing profile to a mesh cluster. If more than two mesh cluster profiles are configured, mesh points use the priority numbers to identify primary and backup profile(s). **NOTE:** The lower the number, the higher the priority. Therefore, the profile with the lowest number is the primary profile. Each profile must use a unique priority value to ensure a deterministic mesh path. Default: 1 for the "default" mesh cluster profile and all user-created cluster profiles. The recovery profile has a priority of 255 (this is not a user-configured profile). The range is 1–16.
Cluster Name	Indicates the mesh cluster name. The name can have a maximum of 32 characters, which is used as the MSSID. When you create a new cluster profile, it is a member of the "Dell-mesh" cluster. **NOTE:** Each mesh cluster profile should have a unique MSSID. Configure a new MSSID before you apply the mesh cluster profile. To view existing mesh cluster profiles, use the command: show ap mesh-cluster-profile. A mesh portal chooses the best cluster profile and provisions it for use. A mesh point can have a maximum of 16 cluster profiles. Default: Mesh cluster named "Dell-mesh."
RF Band	Indicates the band for mesh operation for multiband radios. Select a or g.

Associating a Mesh Cluster Profile to Mesh APs

Use the following procedure to associate a mesh cluster profile to a group of mesh APs or an individual mesh AP using the WebUI. If you configure multiple cluster profiles with different cluster priorities, you manually override the link metric algorithm because the priority takes precedence over the path cost. In this scenario, the mesh portal uses the profile with the highest priority to bring-up the mesh network.

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab.
   - If you selected **AP Group**, click **Edit** by the AP group name to which you want to assign a new mesh cluster profile.
   - If you selected **AP Specific**, click **Edit** by the AP to which you want to assign a new mesh cluster profile
2. Under the Profiles list, expand the **Mesh** menu, then select **Mesh Cluster profile**.
3. In the **Profile Details** window pane, click the **Mesh Cluster profile** drop-down list select **New**.
   - To add an existing mesh cluster profile to the selected AP group, click the **Add a profile** drop-down list and select a new profile name from the list.
   - To create a new mesh cluster profile to the selected AP group, click the **Add a profile** drop-down list and select **NEW**. Enter a name for the new mesh cluster profile.
4. Click the **using priority** drop-down list to select a priority for the mesh cluster profile. The lower the number, the higher the priority.
5. Click **Add** to add the mesh cluster profile to the AP group.
6. Click **Apply**. The profile name appears in the mesh cluster profile list with your configured settings. If you configure this for the AP group, this profile also becomes the mesh cluster profile used by the mesh portal for your mesh network.

Editing a Mesh Cluster Profile

If you modify any mesh cluster profile setting, you must reprovision your AP. For example, if you change the priority of a cluster profile from 5 to 2, you must reprovision the AP before you can assign priority 5 to another cluster profile. Reprovisioning the AP causes it to automatically reboot. For more information, see "Provisioning Mesh Nodes".

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab.
   - If you selected the **AP Group** tab, click the **Edit** button by the AP group name with the profile you want to edit.
   - If you selected the **AP Specific** tab, click the **Edit** button by the AP with the profile you want to edit.
2. In the **Profiles** list, expand the **Mesh** menu, then select **Mesh Cluster profile**.
3. In the **Profile Details** window pane, click the **Mesh Cluster profile** drop-down list and select the name of the profile you want to edit.
4. Change the desired mesh radio settings as desired. Table 104 describes the parameters you can configure in the mesh high-throughput SSID profile.



A mesh cluster profile configured with **wpa2-psk-aes encryption** must have a defined WPA hexkey or a WPA passphrase (or both). If you have configured one encryption type but not the other, and want switch from a hexkey to a passphrase or vice versa, you must add the new encryption type, click **Apply**, then remove the encryption type you no longer want and click **Apply** again. You cannot delete one encryption type and add a different type in a single step.

5. Click **Apply** to save your changes.

Deleting a Mesh Cluster Profile

You can delete a mesh cluster profile only if no APs or AP groups are associated with that profile.

1. Navigate to the **Configuration > Advanced Services> All Profiles** window.
2. Expand the **Mesh** menu, then select **Mesh Cluster profile**. A list of high-throughput SSID profiles appears in the *Profile Details* window pane.
3. Click the **Delete** button by the name of the profile you want to delete.

Managing Mesh Cluster Profiles in the CLI

You must be in config mode to create, modify or delete a mesh cluster profile using the CLI. Specify an existing mesh cluster profile with the <profile-name> parameter to modify an existing profile, or enter a new name to create an entirely new profile.

Configuration details and any default values for each of these parameters are described in Table 102. If you do not specify a parameter for a new profile, that profile uses the default value for that parameter.

Use the **no** option before any parameter to remove the current value for that parameter and return it to its default setting. Enter **exit** to leave the mesh cluster profile mode.

```
(host)(config) #ap mesh-cluster-profile <profile>
  clone <profile>
  cluster <name>
  no ...
  opmode [opensystem | wpa2-psk-aes]
  rf-band {a | g}
  wpa-hexkey <wpa-hexkey>
  wpa-passphrase <wpa-passphrase>
```

The following examples create and configure the mesh cluster profiles **cluster1** and **cluster2**.

```
(host)(config) #ap mesh-cluster-profile cluster1
  cluster corporate
  opmode wpa2-psk-aes
  wpa-passphrase mesh_123
  rf-band a

(host)(config) #ap mesh-cluster-profile cluster2
  cluster corporate
  opmode wpa2-psk-aes
  wpa-passphrase mesh_123rf-band a
```

You can also create a new mesh radio profile by copying the settings of an existing profile using the clone parameter. Using the **clone** command to create a new profile makes it easier to keep constant attributes in common within multiple profiles.

```
(host)(config) #ap mesh-cluster-profile <profile-name> clone <source-profile-name>
```

Viewing Mesh Cluster Profile Settings

To view a complete list of mesh cluster profiles and their status:

```
(host)(config) #show mesh-cluster-profile
```

To view the settings of a specific mesh cluster profile:

```
(host)(config) #show ap mesh-cluster-profile <profile-name>
```

Associating Mesh Cluster Profiles

The following commands associate a mesh cluster profile to an AP group or an individual AP. For deployments with multiple mesh clusters, you must also configure also the profile's priority. Remember, the lower the priority number, the high the priority. The mesh cluster priority determines the order by which the mesh

cluster profiles are used. This allows you, rather than the link metric algorithm, to control the network topology by defining the cluster profiles to use if one becomes unavailable.

To associate a mesh cluster profile to an AP group in a single-cluster deployment:

```
(host)(config) #ap-group <group> mesh-cluster-profile <profile-name>
```

To associate a mesh cluster profile to an individual AP in a single-cluster deployment:

```
(host)(config) #ap-name <name> mesh-cluster-profile <profile-name>
```

To associate a mesh cluster profile to an AP group in a multiple-cluster deployment:

```
(host)(config) #ap-group <group> mesh-cluster-profile <profile-name> priority <priority>
```

To associate a mesh cluster profile to an individual AP in a multiple-cluster deployment, use the command

```
(host)(config) #ap-name <name>
  mesh-cluster-profile <profile-name> priority <priority>
```

Example:

```
(host)(config) #ap-group group1
  mesh-cluster-profile cluster1 priority 5
  mesh-cluster-profile cluster2 priority 10
(host)(config) #ap-group2
  mesh-cluster-profile cluster1 priority 10
  mesh-cluster-profile cluster2 priority 5
  mesh-radio-profile channel2
```

Excluding a Mesh Cluster Profile from a Mesh Node

To exclude a specific mesh cluster profile from an AP:

```
(host)(config) #ap-name <name> exclude-mesh-cluster-profile-ap <profile-name>
```

Deleting a Mesh Cluster Profile

If no AP or is using a mesh cluster profile, you can delete that profile using the **no** parameter:

```
(host)(config) #no ap mesh-cluster-profile <profile-name>
```

Creating and Editing Mesh Radio Profiles

The mesh radio profile determines many of the settings used by mesh nodes to establish mesh links and the path to the mesh portal, including the maximum number of children a mesh node can accept, and transmit rates for the 802.11a and 802.11g radios. The attributes of the mesh radio profile are applied to a mesh point upon receiving its configuration from the controller. You can configure multiple radio profiles; however, you select and deploy only *one* radio profile per AP group. Radio profiles, including the "default" profile, are not active until you provision your APs for mesh.

If you modify a currently provisioned and running radio profile, your changes take effect immediately. You do not need to reboot the controller or the AP to apply the changes.

Managing Mesh Radio Profiles in the WebUI

Use the following procedures to define and manage mesh radio profiles using the WebUI.

Creating or Editing a Mesh Radio Profile

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab.
   - If you selected the **AP Group** tab, click **Edit** by the AP group name for which you want to configure the new mesh radio profile.

- If you selected the **AP Specific** tab, click **Edit** by the AP for which you want to create the mesh radio profile.

2. In the **Profiles** list, expand the **Mesh** menu, then select **Mesh radio profile**.
3. The procedure to create a new mesh profile varies slightly from the procedure to edit an existing profile.
   - To create a new mesh profile: in the **Profile Details** window pane, click the **Mesh radio profile** drop-down list and select **New**. Enter a new mesh radio profile name in the field to the right of the drop-down list.
   - To edit an existing mesh profile: in the **Profile Details** window pane, click the **Mesh radio profile** drop-down list and select the name of the profile you want to edit.
4. Configure your desired mesh radio settings.

   Mesh Radio profile configuration settings are divided into two tabs, **Basic** and **Advanced**. The **Basic** tab displays only those configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. If you change a setting on one tab, then click and display the other tab without saving your configuration, that setting revert sto its previous value. The basic and advanced profile settings are described in Table 103.

Table 103: *Mesh Radio Profile Configuration Parameters*

Parameter	Description
Basic Mesh Radio Settings	
Link Threshold	Use this setting to optimize operation of the link metric algorithm. Indicates the minimal RSSI value. If the RSSI value is below this threshold, the link may be considered a sub-threshold link. A sub-threshold link is one whose average RSSI value falls below the configured link threshold. If this occurs, the mesh node may try to find a better link on the same channel and cluster (only neighbors on the same channel are considered). Default: 12. The supported threshold is hardware dependent, with a practical range of 10–90.
Advanced Mesh Radio Settings	
802.11a Transmit Rates	Indicates the transmit rates for the 802.11a radio. The AP attempts to use the highest transmission rate to establish a mesh link. If a rate is unavailable, the AP goes through the list and uses the next highest rate. To modify transmit rates, do one of the following: • In the WebUI, deselect (uncheck) a specific rate box to use fewer rates when establishing a mesh link. • In the CLI, enter the specific rates to use. Default: All transmission rates are selected and used. If you do not select 802.11a or 802.11g transmit rates, all rates are selected by default when you click **Apply**.
802.11g Transmit Rates	Indicates the transmit rates for the 802.11g radio. The AP attempts to use the highest transmission rate to establish a mesh link. If a rate is unavailable, the AP goes through the list and uses the next highest rate. To modify transmit rates, do one of the following: • In the WebUI, deselect (uncheck) a specific rate box to use fewer rates when establishing a mesh link. • In the CLI, enter the specific rates to use. Default: All transmission rates are selected and used. If you do not select 802.11a or 802.11g transmit rates, all rates are selected by default when you click **Apply**.

Parameter	Description
Allowed VLANs on Mesh Link	List the VLAN ID numbers of VLANs allowed on the mesh link.
BC/MC Rate Optimization	Broadcast/Multicast Rate Optimization dynamically selects the rate for sending broadcast/multicast frames on any BSS. This feature determines the optimal rate for sending broadcast and multicast frames based on the lowest of the unicast rates across all associated clients. When you enable the Multicast Rate Optimization feature, the controller scans the list of all associated stations in that BSS and finds the lowest transmission rate as indicated by the rate adaptation state for each station. If there are no associated stations in the BSS, it selects the lowest configured rate as the transmission rate for broadcast and multicast frames. This feature is enabled by default. Multicast Rate Optimization applies to broadcast and multicast frames only. 802.11 management frames are not affected by this feature and are transmitted at the lowest configured rate. When enabled, this setting dynamically adjusts the multicast rate to that of the slowest connected mesh child. Multicast frames are not sent if there are no mesh children. **NOTE:** This feature should only be enabled on a BSS where all associated stations are sending or receiving unicast data. If there is no unicast data to or from a particular station, then the rate adaptation state may not accurately reflect the current sustainable transmission rate for that station. This could result in a higher packet error rate for broadcast/multicast packets at that station. Configuring the **Video Multicast Rate Optimization** parameter overrides the configuration of **BC/MC Rate Optimization** parameter for VI-tagged multicast traffic. Multicast traffic that is not VI-tagged behaves the same with BC/MC as before. If multicast rate is not set, all traffic behaves the same. Default: Enabled.
Heartbeat threshold	Indicates the maximum number of heartbeat messages that can be lost between neighboring mesh nodes. Default: 10 missed heartbeats. Range: 1–255.
Maximum Children	Indicates the maximum number of children a mesh node can accept. Default: 64 children. Range: 1–64
Maximum Hop Count	Indicates the maximum hop count from the mesh portal. Default: 8 hops. Range: 1–32
Mesh Private VLAN	A VLAN ID for control traffic between an remote mesh portal and mesh nodes. This VLAN ID must not be used for user traffic. Range: 0–4094. Default: 0 (disabled). For further information on configuring a remote mesh portal, see Configuring Remote Mesh Portals (RMPs) on page 567
Mesh Survivability	This feature is currently not supported and should only be enabled under the supervision of Dell support.
Metric algorithm	This parameter specifies the algorithm used by a mesh node to select its parent. Use this setting to optimize operation of the link metric algorithm. Available options are: • best-link-rssi: Selects the parent with the strongest RSSI, regardless of the number of children a potential parent has. • distributed-tree-rssi: selects the parent based on link-RSSI and node cost based on the number of children. This option evenly distributes the mesh points over high quality uplinks. Low quality uplinks are selected as a last resort. Default: distributed-tree-rssi. It is recommended to use the default value.

Parameter	Description
Rate Optimization for delivering EAPOL frames and mesh echoes	When you enable this parameter, EAPOL frames, mesh echo requests and echo responses are sent at a lower rate.
Reselection mode	Use this setting to optimize operation of the link metric algorithm. Available options are: • reselect-anytime • reselect-never • startup-subthreshold • subthreshold-only For complete information on reselection mode options, see Mesh Radio Profiles on page 538
Retry Limit	Indicates the number of times a mesh node can re-send a packet. Default: 4 times. Range: 1–15
RTS Threshold	Defines the packet size sent by mesh nodes. Mesh nodes transmitting frames larger than this threshold must issue request to send (RTS) and wait for other mesh nodes to respond with clear to send (CTS) to begin transmission. This helps prevent mid-air collisions. Default: 2,333 bytes. Range: 256– 2,346.

5. Click **Apply**. The profile name appears in the Mesh Radio Profile list with your configured settings.

   If you configure this for the AP group, this profile also becomes the selected radio profile used by the mesh portal for your mesh network.

Assigning a Mesh Radio Profile to a Mesh AP or AP Group

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab.
   - If you selected **AP Group**, click the **Edit** button by the AP group to which you want to assign a new mesh radio profile.
   - If you selected **AP Specific,** click the **Edit** button by the AP to which you want to assign a new mesh radio profile.
2. Under the **Profiles** list, expand the **Mesh** menu, then select **Mesh radio profile**.
3. In the **Profile Details** window pane, click the **Mesh radio profile** drop-down list and select the desired mesh radio profile from the list.
4. Click **Apply**. The profile name appears in the Mesh Radio Profile list with your configured settings. If you configure this for the AP group, this profile also becomes the selected radio profile used by the mesh portal for your mesh network.
5. Click the **Delete** button by the name of the profile you want to delete.

Managing Mesh Radio Profiles in the CLI

You must be in config mode to create, modify or delete a mesh radio profile using the CLI. Specify an existing mesh profile with the <profile-name> parameter to modify an existing profile, or enter a new name to create an entirely new profile.

Creating or Modifying a Mesh Radio Profile

Configuration details and any default values for each of these parameters are described in Table 103. If you do not specify a parameter for a new profile, that profile uses the default value for that parameter. Put the **no** option before any parameter to remove the current value for that parameter and return it to its default setting. Enter **exit** to leave the mesh radio profile mode.

```
(host)(config) #ap mesh-radio-profile <profile-name>
  a-tx-rates
  allowed-vlans
  children <children>
  clone <source-profile-name>
  eapol-rate-opt
  g-tx-rates [1|2|5|6|9|11|12|18|24|36|48|54]
  heartbeat-threshold <count>
  hop-count <hop-count>
  link-threshold <count>
  max-retries <max-retries>
  mesh-ht-ssid-profile
  mesh-mcast-opt
  mesh-survivability
  metric-algorithm {best-link-rssi|distributed-tree-rssi}
  mpv <vlan-id>
  no
  reselection-mode
  rts-threshold <rts-threshold>
```

You can also create a new mesh radio profile by copying the settings of an existing profile using the clone parameter. Using the `clone` command to create a new profile makes it easier to keep constant attributes in common within multiple profiles.

```
(host)(config) #ap mesh-radio-profile <profile-name> clone <source-profile-name>
```

Assigning a Mesh Radio Profile to a Mesh AP or AP Group

To associate a mesh radio profile with an AP group, use the following commands. When you add the mesh cluster profile to the AP group, you must also define the cluster priority.

```
(host)(config) #ap-group <group>
  mesh-radio-profile <profile-name> priority <priority>
```

To associate a mesh radio profile with an individual AP:

```
(host)(config) #ap-name <name>
  mesh-radio-profile <profile-name> priority <priority>
```

The following examples assign the mesh cluster profiles **cluster1** and **cluster2** to two different AP groups. In the AP group **group1, cluster1** has a priority of 5, and **cluster2** has a priority of 10, so **cluster1** has the higher priority. In the AP group **group2**, **cluster1** has a priority of 10, and **cluster2** has a priority of 5, so **cluster5** has the higher priority.

```
(host)(config) #group2—cluster1 has a priority of 10, and cluster2 has a priority of 5.
(host)(config) #ap-group group1
      mesh-cluster-profile cluster1 priority 5
      mesh-cluster-profile cluster2 priority 10

(host)(config) #ap-group2
  mesh-cluster-profile cluster1 priority 10
  mesh-cluster-profile cluster2 priority 5
```

Deleting Mesh Radio Profiles

You can delete a mesh radio profile only if no other APs or AP groups use that profile.

To delete a mesh radio profile using the WebUI:

1. Navigate to the **Configuration > Advanced Services> All Profiles** window.
2. Expand the **Mesh** menu, then select **Mesh radio profile**. A list of mesh radio profiles appears in the **Profile Details** window pane.
3. Click **Delete** by the name of the profile you want to delete.

The following command deletes a radio profile via the command-line interface.

```
(host)(config)no ap mesh-radio-profile <profile-name>
```

Creating and Editing Mesh High-Throughput SSID Profiles

The mesh high-throughput SSID profile defines settings unique to 802.11n-capable, high-throughput APs. If none of the APs in your mesh deployment are 802.11n-capable APs, you do not need to configure a high-throughput SSID profile. If you modify a currently provisioned and running high-throughput SSID profile, your changes take effect immediately. You do not reboot the controller or the AP.

Managing Mesh High-Throughput SSID Profiles in the WebUI

Use the following procedures to manage your high-throughput SSID profiles using the WebUI.

Creating a Profile

To create a high-throughput SSID profile:

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab.
   - If you selected **AP Group**, click **Edit** by the AP group for which you want to create the new high-throughput SSID profile.
   - If you selected **AP Specific**, click **Edit** by the AP for which you want to create the new high-throughput SSID profile.
2. In the **Profiles** list, expand the **Mesh** menu, then select **Mesh High-throughput SSID profile**.
3. In the **Profile Details** window pane, click the **Mesh High-throughput SSID profile** drop-down list and select **NEW**.
4. Enter a name for the new profile.
5. Configure the mesh high-throughput SSID parameters described in Table 104. The Mesh High-Throughput SSID Profile configuration settings are divided into two tabs, **Basic** and **Advanced**. The **Basic** tab displays only those configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. If you change a setting on one tab then click and display the other tab without saving your configuration, that setting reverts to its previous value.
6. Click **Apply**. The profile name appears in the Mesh High-throughput SSID Profile list with your configured settings.

Table 104: *Mesh High-Throughput SSID Profile Configuration Parameters*

Parameter	Description
Basic Mesh High-Throughput SSID Profile Settings	
40 MHz channel usage	Enable or disable the use of 40 MHz channels. Default: enabled
High-throughput Enable (SSID)	Enable or disable high-throughput (802.11n) features on the SSID. Default: enabled
Explicit Transmit Beamforming	Enable/Disable use of Explicit Transmit Beamforming. (For W-AP130 Series only) If this parameter is disabled, the other transmit beamforming configuration settings have no effect.
Transmit Beamforming Compressed Steering	When enabled, the AP can use explicit compressed feedback from clients to obtain a steering matrix. (For W-AP130 Series APs only.) Default: enabled
Transmit Beamforming non Compressed Steering	When enabled, the AP can use explicit noncompressed feedback from clients to obtain a steering matrix. (For W-AP130 Series only) Default: enabled
Transmit Beamforming delayed feedback support	Enable/Disable delayed feedback/report support in Transmit Beamforming. (For W-AP130 Series only) Default: enabled
Transmit Beamforming immediate feedback support	Enable/Disable immediate feedback/report support in Transmit Beamforming. (For W-AP130 Series only) Default: enabled
Transmit Beamforming Sounding Interval	Time interval in seconds between updates of Transmit Beamforming channel estimation. (For W-AP130 Series only) The supported range is 1-65335 seconds, and the default is 1800 seconds.
Advanced Mesh High-Throughput SSID Profile Settings	
Temporal Diversity Enable	When a client is not responding to 802.11 packets, the AP will launch two hardware retries. If you enable this option and hardware retries are not successful, then the AP will launch then software retries.
BA AMSDU Enable	Enable/Disable Receive AMSDU in BA negotiation.
Legacy stations	Allow or disallow associations from legacy (non-HT) stations. By default, this parameter is enabled (legacy stations are allowed).
Low-density Parity Check	If enabled, the AP advertises Low-density Parity Check (LDPC) support LDPC improves data transmission over radio channels with high levels of background noise. (For W-AP130 Series only)
Maximum number of spatial streams usable for STBC reception	Controls the maximum number of spatial streams usable for STBC reception. 0 disables STBC reception, 1 uses STBC for MCS 0-7. Higher MCS values are not supported. (Supported on the W-AP90 series, W-AP130 Series, W-AP68, W-AP175 and W-AP105 only. The configured value adjusts based on AP capabilities.) If transmit beamforming is enabled, STBC is disabled for disabled for beamformed frames.

Parameter	Description
Maximum number of spatial streams usable for STBC transmission.	Controls the maximum number of spatial streams usable for STBC transmission. 0 disables STBC transmission, 1 uses STBC for MCS 0-7. Higher MCS values are not supported. (Supported on W-AP90 series, W-AP175, W-AP130 Series and W-AP105 only. The configured value adjusts based on AP capabilities.) If you enable transmit beamforming, STBC is disabled for disabled for beamformed frames.
MPDU Aggregation	Enable or disable MAC protocol data unit (MPDU) aggregation. High-throughput APs are able to send aggregated MAC protocol data units (MDPUs), which allow an AP to receive a single block acknowledgment instead of multiple ACK signals. This option, which is enabled by default, reduces network traffic overhead by effectively eliminating the need to initiate a new transfer for every MPDU.
Max received A-MPDU size	Maximum size of a received aggregate MPDU, in bytes. Allowed values: 8191, 16383, 32767, 65535.
Max transmitted A-MPDU size	Maximum size of a transmitted aggregate MPDU, in bytes. Range: 1576–65535
Min MPDU start spacing	Minimum time between the start of adjacent MPDUs within an aggregate MPDU, in microseconds. Allowed values: 0 (No restriction on MDPU start spacing), .25 µsec, .5 µsec, 1 µsec, 2 µsec, 4 µsec.

Parameter	Description
Short guard interval in 20 MHz mode	Enable or disable use of short (400ns) guard interval in 20 MHz mode. This parameter is enabled by default. A guard interval is a period of time between transmissions that allows reflections from the previous data transmission to settle before an AP transmits data again. An AP identifies any signal content received inside this interval as unwanted inter-symbol interference, and rejects that data. The 802.11n standard specifies two guard intervals: 400ns (short) and 800ns (long). Enabling a short guard interval can decrease network overhead by reducing unnecessary idle time on each AP. Some outdoor deployments, may, however require a longer guard interval. If the short guard interval does not allow enough time for reflections to settle in your mesh deployment, inter-symbol interference values may increase and degrade throughput.
Short guard interval in 40 MHz mode	Enable or disable use of short (400ns) guard interval in 40 MHz mode. This parameter is enabled by default. A guard interval is a period of time between transmissions that allows reflections from the previous data transmission to settle before an AP transmits data again. An AP identifies any signal content received inside this interval as unwanted inter-symbol interference, and rejects that data. The 802.11n standard specifies two guard intervals: 400ns (short) and 800ns (long). Enabling a short guard interval can decrease network overhead by reducing unnecessary idle time on each AP. Some outdoor deployments, may, however require a longer guard interval. If the short guard interval does not allow enough time for reflections to settle in your mesh deployment, inter-symbol interference values may increase and degrade throughput.
Supported MCS set	A list of Modulation Coding Scheme (MCS) values or ranges of values to be supported on this SSID. The MCS you choose determines the channel width (20MHz vs. 40MHz) and the number of spatial streams used by the mesh node. The default value is 1–23; the complete set of supported values. To specify a smaller range of values, enter a hyphen between the lower and upper values. To specify a series of different values, separate each value with a comma. Examples: 2–10 1,3,6,9,12 Range: 0–23.

Assigning a Profile to an AP Group

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab.
   - If you selected **AP Group**, click **Edit** by the AP group name to which you want to assign a new high-throughput SSID profile.
   - If you selected **AP Specific**, click **Edit** by the AP to which you want to assign a new high-throughput SSID profile
2. Under the **Profiles** list, expand the **Mesh** menu, then select **Mesh High-throughput SSID profile**.
3. In the **Profile Details** window pane, click the **Mesh High-throughput SSID profile** drop-down list and select the desired profile from the list.
4. Click **Apply**. The profile name appears in the Mesh High-throughput SSID Profile list with your configured settings. If you configure this for the AP group, this profile also becomes the selected high-throughput SSID profile used by the mesh portal for your mesh network.

Editing a Profile

1. Navigate to the **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab.

- If you selected the **AP Group** tab, click **Edit** by the AP group name with the profile you want to edit.
- If you selected the **AP Specific** tab, click **Edit** by the AP with the profile you want to edit.

2. In the **Profiles** list, expand the **Mesh** menu, then select **Mesh High-throughput SSID profile**.
3. In the **Profile Details** window pane, click the **Mesh High-throughput SSID profile** drop-down list and select the name of the profile you want to edit.
4. Change the settings as desired. Table 104 describes the parameters you can configure in this profile.
5. Click **Apply**.

Deleting a Profile

You can delete a mesh high-throughput SSID profile only if no APs or AP groups are associated with that profile.

1. Navigate to the **Configuration > Advanced Services> All Profiles** window.
2. Expand the **Mesh** menu, then select **Mesh High-throughput SSID profile**. A list of high-throughput SSID profiles appears in the **Profile Details** window pane.
3. Click **Delete** by the name of the profile you want to delete.

Managing Mesh High-Throughput SSID Profiles in the CLI

You must be in config mode to create, modify or delete a mesh high-throughput SSID radio profile using the CLI. Specify an existing high-throughput SSID profile with the <profile-name> parameter to modify an existing profile, or enter a new name to create an entirely new profile.

Creating or Modifying a Profile

Configuration details and any default values for each of these parameters are described in Table 104. If you do not specify a parameter for a new profile, that profile uses the default value for that parameter. Put the **no** option before any parameter to remove the current value for that parameter and return it to its default setting. Enter **exit** to leave the high-throughput radio profile mode

```
(host)(config) #ap mesh-ht-ssid-profile <profile-name>
  40MHz-enable
  clone
  high-throughput-enable
  ldpc
  legacy-stations
  max-rx-a-mpdu-size
  max-tx-a-mpdu-size
  min-mpdu-start-spacing
  mpdu-agg
  no
  short-guard-intvl-20mhz
  short-guard-intvl-40mhz
  stbc-rx-streams
  stbc-tx-streams
  supported-mcs-set
  temporal-diversity
```

You can also create a new mesh high-throughput SSID profile by copying the settings of an existing profile using the clone parameter. Using the `clone` command to create a new profile makes it easier to keep constant attributes in common within multiple profiles.

```
ap mesh-ht-ssid-profile <profile-name> clone <source-profile-name>
```

Assigning a Profile to an AP Group

To associate a mesh high-throughput SSID profile with an AP group:

```
(host)(config) #ap-group <group> mesh-ht-ssid-profile <profile-name>
```

To associate a mesh radio profile with an individual AP:

```
(host)(config) #ap-name <name> mesh-ht-ssid-profile <profile-name>
```

Viewing High-throughput SSID Settings

To view a complete list of high-throughput profiles and their status:

```
(host)(config) #show ap mesh-ht-ssid-profile
```

To view the settings of a specific high-throughput profile:

```
(host)(config) #show ap mesh-ht-ssid-profile <profile-name>
```

Deleting a Profile

If no AP or AP group is using a mesh high-throughput SSID profile, you can delete that profile using the **no** parameter:

```
(host)(config) no ap mesh-ht-ssid-profile <profile-name>
```

Configuring Ethernet Ports for Mesh

If you use mesh to join multiple Ethernet LANs, configure and enable bridging on the mesh point Ethernet port This section describes how to configure Ethernet ports for bridging or secure jack operation using the wired AP profile. The wired AP profile controls the configuration of the Ethernet port(s) on your AP.



Mesh nodes only support bridge mode and tunnel mode on their wired ports (enet0 or enet1). Split tunnel mode is not supported. Use bridge mode to configure bridging on the mesh point Ethernet port. Use tunnel mode to configure secure jack operation on the mesh node Ethernet port.

When configuring the Ethernet ports on dual-port APs, note the following requirements for the AP configures as a mesh portal:

- Connect enet0 to the controller to obtain an IP address. The wired AP profile controls enet1.
- Only enet1 supports secure jack operation.

Configuring Bridging on the Ethernet Port

Use the following procedure to configure bridging on the Ethernet port via the WebUI.

1. Navigate to the **Configuration > Wireless > AP Configuration > AP Group** window.
2. Click the **Edit** button by the AP group name with the wired ap profile you want to edit.
3. Under the Profiles list, expand the **AP menu**, then select **Wired AP profile**. The settings for the currently selected wired AP profile appear.

   You can use a different wired AP profile by selecting a profile from the **Wired AP profile** drop-down list.
4. Under Profile Details, do the following:
   a. Select the **Wired AP enable** check box. This option is not selected by default.
   b. From the **Forward mode** drop-down list, select **bridge**.
   c. Optionally, from the **Switchport mode** drop-down list, select **access or trunk**. These options only apply to bridge mode configurations.
      - Access mode forwards untagged packets received on the port to the controller and they appear on the configured access mode VLAN. Tagged packets are dropped. All packets received from the controller and sent via this port are untagged. Define the access mode VLAN in the **Access mode VLAN** field.
      - Trunk mode contains a list of allowed VLANs. Any packet received on the port that is tagged with an allowed VLAN is forwarded to the controller. Untagged packets are forwarded to the controller on the

configured Native VLAN. Packets received from the controller and sent out the port remain tagged unless the tag value in the packet is the Native VLAN, in which case the tag is removed. Define the Native VLAN in the **Trunk mode native VLAN** field and the other allowed VLANs in the **Trunk mode allowed VLANs** field.

d. Optionally, select **Trusted** to configure this as a trusted port.

5. Click **Apply**.

Use the following commands to configure Ethernet port bridging via the CLI.

```
(host)(config) #ap wired-ap-profile <profile>
   forward-mode bridge
   wired-ap-enable
```

Optionally, you can configure the following wired AP profile settings:

```
(host)(config) #ap wired-ap-profile <profile>
   switchport mode {access | trunk}
   switchport access vlan <vlan>
   switchport trunk native vlan <vlan>
   switchport trunk allowed vlan <vlan>
   trusted
```

Configuring Ethernet Ports for Secure Jack Operation

You can configure the Ethernet port(s) on mesh nodes to operate in tunnel mode. Known as secure jack operation for mesh, this configuration allows Ethernet frames coming into the specified wired interface to be generic routing encapsulation (GRE) tunneled to the controller. Likewise, Ethernet frames coming from the tunnel are bridged to the corresponding wired interface. This allows an Ethernet port on the mesh node to appear as an Ethernet port on the controller separated by one or more Layer-3 domains. You can also enable VLAN tagging.

Unlike secure jack on non-mesh APs, any mesh node configured for secure jack uses the mesh link, rather than enet0, to tunnel the frame to the controller.

When configuring mesh Ethernet ports for secure jack operation, note the following guidelines:

- Mesh points support secure jack on enet0 and enet1.
- Mesh portals only support secure jack on enet1. This function is only applicable to Dell APs that support a second Ethernet port and mesh, such as the W-AP130 Series.

You configure secure jack operation in the wired AP profile.



The parameters in the wired AP profile only apply to the wired AP interface to which they are applied. Two wired interfaces can have different parameter values.

In the WebUI

Use the following procedure to configure secure jack operation using the WebUI.

1. Navigate to the **Configuration > Wireless > AP Configuration > AP Group** window.
2. Click **Edit** by the AP group with the wired AP profile you want to edit.
3. Under the **Profiles** list, expand the **AP** menu**,** then select **Wired AP profile**. The settings for the currently selected wired AP profile appear.

   You can use a different wired AP profile by selecting a profile from the **Wired AP profile** drop-down list.
4. In the **Profile Details** window pane, do the following:
   a. Select the **Wired AP enable** check box. This option is not selected by default.
   b. From the **Forward mode** drop-down list, select **tunnel**.
   c. Optionally, select **Trusted** to configure this as a trusted port.

5. Click **Apply**.

In the CLI

To configure secure jack operation using the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #ap wired-ap-profile <profile>
   forward-mode tunnel
   wired-ap-enable
```

Optionally, you can configure the following wired AP profile settings:

```
(host)(config) #ap wired-ap-profile <profile>
   trusted
```

Extending the Life of a Mesh Network

To prevent your mesh network from going down if you experience a controller failure, modify the following settings in the AP system profile(s) used by mesh nodes to maintain the mesh network until the controller is available:



It is recommended to use the default maximum request retries and bootstrap threshold settings for most mesh networks; however, if you must keep your mesh network alive, you can modify the settings as described in this section. The modified settings are not applicable if mesh portals are directly connected to the controller.

- Maximum request retries: maximum number of times to retry AP-generated requests. The default is 10 times. If you must modify this setting, it is recommended to set a value of 10,000.
- Bootstrap threshold: number of consecutive missed heartbeats before the AP rebootstraps. (Heartbeats are sent once per second.) The default is 9 missed heartbeats. If you must modify this setting, it is recommended to set a value of 5,000.

When the controller comes back online, the affected mesh nodes (mesh portals and mesh points) rebootstrap; however, the mesh link is not affected and continues to be up.

In the WebUI

Use the following procedure to modify the AP system profile via the WebUI.

1. Navigate to the **Configuration > Wireless > AP Configuration > AP Group** window.
2. Click the **Edit** button by the AP group with the AP system profile you want to edit.
3. Under **Profiles** list, expand the **AP** menu, then select **AP system profile**. The settings for the currently selected AP system profile appear in the **Profile Details** window pane.
4. Make the following changes in the **Profile Details** window pane.
   a. Change the **Maximum Request Retries** to 10000.
   b. Change the **Bootstrap threshold** to 5000.
5. Click **Apply**.

In the CLI

To modify the AP system profile via the command-line interface, access the CLI in config mode and issue the following commands:

```
(host)(config) #ap system-profile <profile>
   max-request-retries 10000
   bootstrap-threshold 5000
```

Provisioning Mesh Nodes

Provisioning mesh nodes is similar to thin APs; however, there are some key differences. Thin APs establish a channel to the controller from which they receive the configuration for each radio interface. Mesh nodes, in contrast, get their radio interfaces up and running before making contact with the controller. This requires a minimum set of parameters from the AP group and mesh cluster that enables the mesh node to discover a neighbor to create a mesh link and subsequent channel with the controller. To do this, you must first configure mesh cluster profiles for each mesh node prior to deployment. See “Mesh Radio Profiles” for more information.

On each radio interface, you provision a mode of operation: mesh node or thin AP (access) mode. If you do not specify mesh, the AP operates in thin AP (access) mode. If you configure mesh, the AP is provisioned with a minimum of two mesh cluster profiles: the “default” mesh cluster profile and an emergency read-only recovery profile, as described in the section “Mesh Clusters”. If you create and select multiple mesh cluster profiles, the AP is provisioned with those as well. If you have a dual-radio AP and configure one radio for mesh and the other as a thin AP, each radio is provisioned as configured.

Each radio provisioned in mesh mode can operate in one of two roles: mesh portal or mesh point. You explicitly configure the role, as described in this section. This allows the AP to know whether it uses the mesh link (via the mesh point/mesh portal) or an Ethernet link to establish a connection to the controller.

During the provisioning process, mesh nodes look for a mesh profile that the AP group and AP name is a member of and stores that information in flash. If you have multiple cluster profiles, the mesh portal uses the best profile to bring-up the mesh network. Mesh points in contrast go through the list of mesh cluster profiles in order of priority to decide which profile to use to associate themselves with the network. In addition, when a mesh point is provisioned, the country code is sent to the AP from its AP name or AP group along with the mesh cluster profiles. Mesh nodes also learn the recovery profile, which is automatically generated by the master controller. If the other mesh cluster profiles are unavailable, mesh nodes use the recovery profile to establish a link to the master controller; data forwarding does not take place.



If you create a new mesh cluster profile for an existing deployment, you must re-provision the AP for the new profile to take effect. If you re-provision mesh nodes that are already operating, re-provision the most distant (highest hop count) mesh points first followed by the mesh portals. If you re-provision the mesh portal first, the mesh points may be unable to form a mesh link. Re-provisioning the AP causes it to automatically reboot. This may cause a disruption of service to the network.

Provisioning Caveats

Remember the following when provisioning APs for mesh:

- You must provision the AP before you install it as a mesh node in a mesh deployment. To provision the AP, it must be physically connected to the local network or directly connected to the controller. When connected and powered on, the AP must also be able to obtain an IP address from a DHCP server on the local network or from the controller.
- Make sure the provisioned mesh nodes form a connected mesh network before physically deploying the APs. For more information, see “Verifying the Network”.
- In multi-controller networks, save your mesh cluster configuration before provisioning the mesh nodes. To save your configuration in the WebUI, at the top of any window click **Save Configuration**. To save your configuration in the CLI, use the command **write memory**.
- If the same port on the controller is used to provision APs and provide PoE for mesh nodes, you must stop traffic from passing through that port after you provision the AP. To stop traffic, shut down (disable) the port either by using the CLI command **interface fastethernet <slot>/<port> shutdown**, or by following the procedure below.

  1. Navigate to the **Configuration > Network > Ports** window.

2. Under **Port Selection**, click the port to configure.
3. Under **Configure Selected Port**, deselect (uncheck) **Enable Port**.
4. Make sure **Enable 802.3af Power Over Ethernet** is selected.
5. Click **Apply**.

Provisioning Mesh Nodes

Reprovisioning the AP causes it to automatically reboot. The following procedures describe the process to provision a mesh portal or mesh node via the WebUI or CLI. (The easiest way to provision a mesh node is to use the Provisioning window in the WebUI.) To provision a remote mesh portal, see "Configuring Remote Mesh Portals (RMPs).

In the WebUI

1. Navigate to the **Configuration > Wireless > AP Installation > Provisioning** window. Select the AP to provision for mesh and click **Provision**.
2. In the **Master Discovery** section, set the Master IP address as the controller IP address.
3. In the **IP settings** section, select **Obtain IP Address Using DHCP**.
4. In the **AP List** section, do the following:
   - Configure the Mesh Role:
     - To configure the AP as the mesh portal, select **Mesh Portal**.
     - To configure the AP as a mesh point, select **Mesh Point**.
   - Configure the Outdoor Parameters, if needed. The following parameters are available only if configuring an outdoor AP:
     - Latitude coordinates (degrees, minutes, seconds, north or south)
     - Longitude coordinates (degrees, minutes, seconds, east or west)
     - Altitude (in meters)
     - Antenna bearing (horizontal coverage)
     - Antenna tilt angle (optimum coverage)



The above parameters apply to all outdoor APs, not just outdoor APs configured for mesh.

5. Click **Apply and Reboot**. After the controller reboots, mesh cluster profiles are extracted from the AP group and the AP name.

In the CLI

When you use the command-line interface to reprovision a mesh node, you may also provision other AP settings.

Access the CLI in config mode and issue the following commands:

```
(host)(config) #provision-ap
   read-bootinfo ap-name <name>
   mesh-role {mesh-point|mesh-portal}
   reprovision ap-name <name>
```

If you are provisioning an outdoor AP, you can also configure the following parameters:

```
(host)(config) #provision-ap
   read-bootinfo ap-name <name>
   mesh-role {mesh-point|mesh-portal}
   a-ant-bearing <bearing>
   a-ant-tilt-angle <angle>
```

```
g-ant-bearing <bearing>
g-ant-tilt-angle <angle>
altitude <altitude>
latitude <location>
longitude <location>
reprovision ap-name <name>
```

Verifying Your Mesh Network

To view a list of your Mesh APs via the WebUI, navigate to the one of the following windows:

- **Monitoring > Network > All Mesh Nodes**
- **Monitoring > Controller > Mesh Nodes**

To view mesh APs and the mesh topology tree using the command line interface, access the command-line interface in enable mode and issue the following commands:

- `show ap mesh active`
- `show ap mesh topology`

Verification Checklist

After provisioning the mesh APs, follow the steps below to ensure that the mesh network is up and operating correctly.

- Issue the command **show ap mesh topology** to verify all the mesh APs are up and the topology is as expected. (Wait 10 minutes after startup for the topology to stabilize.)
- Verify each mesh node has the expected RSSI to its neighboring mesh nodes. The mesh topology is updated periodically, so access the command-line interface and issue the command **`show ap mesh neighbors`** for the current status. If the RSSI is low, verify that the tx-power settings in the mesh node's 802.11a/802.11g radio profiles are correct, or, if ARM is used, verify the correct minimum tx-power setting.
- Issue the command **show ap mesh debug provisioned-clusters** to verify that the mesh clusters are correctly defined and provisioned (with encryption if desired). Issue the **show running-config | include recovery** command to verify that the cluster's recovery profile matches the controller's recovery profile.
- Verify antenna provisioning by issuing the **show ap provisioning** command and verify installation parameters for non-default installations (that is, standard indoor APs deployed outside, or outdoor APs deployed inside). Ensure all APs use the same channel list by issuing the **show ap allowed-channels** command.
- *If the mesh-radio is to be reserved exclusively for mesh backhaul traffic,* issue the command **show ap profile-usage** to identify the radio's 802.11a or 802.11g radio profile, then issue the command **show rf dot11a-radio-profile <profile>** or **show rf dot11g-radio-profile <profile>** to verify the radio is disabled in the profile. Next, use the **show ap bss-table** command to that verify no access Virtual APs are up on the mesh radio.

CLI Examples

Use the **show ap mesh active** command to verify all nodes are present and that EIRP is correct:

```
(host) #show ap mesh active
Mesh Cluster Name: meshprofile1
-------------------------------
Name    Group    IP Address      BSSID             Band/Ch/EIRP/MaxEIRP MTU Enet 0/1 Mesh Role
----    -----    ----------       -----            -------------------- --- -------- ---------
mp1      mp1     10.3.148.245    00:1a:1e:85:c0:30  802.11a/157/19/36 Off/Off Point
mp2      mp2     10.3.148.250    00:1a:1e:88:11:f0  802.11a/157/19/36 Bridge/Bridge Point
mp3      mp3     10.3.148.253    00:1a:1e:88:01:f0  802.11a/157/19/36 Bridge/Bridge Point
mpp      mpp125  10.3.148.252    00:1a:1e:88:05:50  802.11a/157/19/36 1578 -/Bridge Portal
```

```
Parent #Children        AP Type         Uptime
------ ---------        -------         ------
mp3     0               125             13d:2h:25m:19s
mpp     1               125             14d:21h:23m:49s
mp2     1               125             14d:21h:14m:55s
-       1               125             14d:19h:5m:3s
```

Use the **show ap mesh topology** command to verify the cluster topology, RSSI in presence of network traffic, and Tx and Rx rates.

```
(host) #show ap mesh topology

Mesh Cluster Name: sw-ad-GB32
-----------------------------
Name Mesh Role    Parent  Path Cost  Node Cost  Link Cost  Hop Count  RSSI  Rate Tx/Rx
---- ---------    ------  ---------  ---------  ---------  ---------  ----  ----------  Last
Update  Uplink Age  #Children
----------  ----------  ---------

ad-ap Point (N)    mp3     2          0          0          1          61    300/270     6m:12s
     3h:8m:7s    0

msc-1 Point        mp3     2          0  0          1          64    54/54       6m:36s
2h:48m:12s  0

Total APs :2
(R): Recovery AP. (N): 11N Enabled. For Portals 'Uplink Age' equals uptime.
```

Issue the command **show ap mesh neighbors ap-name <name>** to verify visibility of other mesh nodes is as expected:(host) #show ap mesh neighbors ap-name portal

```
Neighbor list
-------------
MAC                Portal             Channel  Age  Hops  Cost   Relation      Flags  RSSI
Rate Tx/Rx
---                ------             -------  ---  ----  ----   --------      -----  ----  --
--------
00:0b:86:e8:09:d1  00:1a:1e:88:01:f0  157      0    1     11.00  C 3h:15m:42s  -      65
54/54
00:1a:1e:88:02:91  00:1a:1e:88:01:f0  157      0    1     4.00   C 3h:35m:30s  HL     59
300/300
00:0b:86:9b:27:78  Yes                157      0    0     12.00  N 3h:22m:46s  -      26    -

00:0b:86:e8:09:d0  00:1a:1e:88:01:f0  157      0    1     11.00  N 3h:15m:36s  -      65    -

00:1a:1e:88:02:90  00:1a:1e:88:01:f0  157+     0    1     2.00   N 3h:35m:6s   HL     59    -


A-Req  A-Resp  A-Fail  HT-Details       Cluster ID
-----  ------  ------  ----------       ----------
1      1       0       Unsupported      sw-ad-GB32
1      1       0       HT-40MHzsgi-2ss  sw-ad-GB322
0      0       0       Unsupported      mc1
0      0       0       Unsupported      sw-ad-GB32
0      0       0       HT-40MHzsgi-2ss  sw-ad-GB32

Total count: 5, Children: 2
```

Configuring Remote Mesh Portals (RMPs)

The following steps describe the procedure to configure a Remote Mesh portal using the WebUI and CLI interfaces.

Creating a Remote Mesh Portal In the WebUI

A remote mesh portal must be provisioned as both a remote access point and a mesh portal. For instructions on provisioning the remote mesh portal as a remote access point, see Configuring the Secure Remote Access Point Service on page 631.

Wired ports on remote mesh portals can be configured in either bridge or split-tunnel forwarding mode. However, there are limitations to the forwarding modes that can be used by other mesh node types. Do not use bridge or split-tunnel forwarding mode for wired ports on mesh points. Virtual APs on remote mesh portals and remote mesh points also do not support bridge or split-tunnel forwarding mode.



A remote mesh portal does not support bridge mode Virtual APs or offline Virtual APs.

Step 1: Provision the AP

1. Navigate to the **Configuration > Wireless > AP Installation > Provisioning** window.
2. Select the AP to provision as a remote mesh portal and click **Provision**. The **Provisioning** window appears.
3. In the **Authentication** section, select the **Remote AP** radio button.
4. In the **Remote AP Authentication Method** section of this window, select either **Pre-shared Key** or **Certificate**. If you selected **Pre-Shared Key**, enter and confirm the Internet Key Exchange Pre-Shared Key (IKE PSK).
5. In the **Master Discovery** section, set the Master IP address as the controller IP address.
6. In the **IP settings** section, select **Obtain IP Address Using DHCP**.
7. In the **AP List** section at the bottom of the window, click the **Mesh Role** drop-down list and select **Remote Mesh Portal**.

Step 2: Define the Mesh Private VLAN in the Mesh Radio Profile

Follow the procedure below to choose a new, non-zero tag value for the mesh private VLAN. Make sure that the mesh private VLAN so that it does not conflict with any local tags assigned in the mesh network. once configured, all Mesh Points come up in that Mesh Private Vlan. This mesh private VLAN must not be used as a VLAN for any other virtual AP.

1. Edit the mesh radio profile for the remote mesh portal according to the procedure described inCreating or Editing a Mesh Radio Profile on page 551 .
2. Set the **Mesh Private VLAN** parameter to define a VLAN ID (0-4094) for control traffic between an remote mesh point and mesh nodes.
3. Click **Apply** to save your changes.

Next, assign the remote mesh points with the same mesh cluster profile, 802.11a and 802.11g RF management profiles, and mesh radio profile as the remote mesh portal. If you have defined an AP group for all your remote mesh points, you can just assign the required profiles to the remote mesh point AP group. Otherwise, you must assign the required profiles to each individual remote AP.

Step 3: Assign the Mesh Radio Profile to a Remote Mesh AP

Follow the procedures described in Assigning a Mesh Radio Profile to a Mesh AP or AP Group on page 554

Step 4: Assign an RF Management Profile to a Remote Mesh AP

Follow the procedures described in Assigning an 802.11a/802.11g Profile to an AP or AP Group on page 515 to assign an 802.11a or 802.1g RF management profile to the remote mesh AP.

Step 5: Assign a Mesh Cluster Profile

Follow the procedures described in Configuring Mesh Cluster Profiles on page 547 to assign a mesh cluster profile to the remote mesh AP.



If you configure multiple cluster profiles with different cluster priorities, you manually override the link metric algorithm because the priority takes precedence over the path cost. In this scenario, the mesh portal uses the profile with the highest priority to bring-up the mesh network.

Step 6: Configuring a DHCP Pool

In this next step, you must configure a DHCP pool where the DHCP server is on the subnet associated with mesh private VLAN. Mesh points get their IP address from this subnet pool. To complete this task, refer to the procedure described in “Configuring the DHCP Server on the Remote AP”.

Step 7: Configuring the VLAN ID of the Virtual AP Profile

Follow the procedure described in SSID Profiles on page 414 to configure the VLAN ID of the remote mesh AP's SSID profile. The VLAN of this Virtual AP must have the same VLAN ID as the mesh private VLAN.

Provisioning a Remote Mesh Portal In the CLI

Reprovisioning the AP causes it to automatically reboot. When you use the CLI to reprovision a mesh node, you may also provision other AP settings.

```
(host)(config) #provision-ap
   read-bootinfo ap-name <name>
   mesh-role remote-mesh-portal
   reprovision ap-name <name>
```

A single controller at the core of a network can represent a single point of failure. ArubaOS high availability and Virtual Router Redundancy Protocol (VRRP) redundancy features allows network administrators to significantly reduce network downtime and client traffic disruption during network upgrades or unexpected failures.

High Availability

When you enable High Availability WLAN redundancy solution, campus APs that lose contact with their active controller do not need to rebootstrap when they fail over to the standby controller, significantly reducing AP downtime. APs using the High Availability features regularly communicate with the standby controller, so the standby controller has only a light workload to process in the event of an AP failover. This results in very rapid failover times, and a shorter client reconnect period. Therefore, High Availability is usually preferable to other redundancy solutions (like a backup-LMS) that can put a heavy load on the backup controller during failover, resulting in slower failover performance.

High Availability supports failover for campus APs using tunnel, decrypt-tunnel, or bridge forwarding modes. It does not support failover for remote APs.



AP Fast Failover on bridge forwarding mode virtual AP is supported on W-7200 Series controllers only.

Pre-Deployment Information

For information to help you plan your high availability solution, refer to the following sections of this document:

- High Availability Deployment Models on page 571
- High Availability Extended Controller Capacity on page 574
- Client State Synchronization on page 573
- High Availability Inter-Controller Heartbeats on page 574

Configuration Procedures

For more information on configuring the high availability feature, refer to the following sections of this document:

- Configuring High Availability on page 576
- Migrating from VRRP or Backup-LMS Redundancy on page 578

VRRP-Based Redundancy

The Virtual Router Redundancy Protocol (VRRP) is used to create various redundancy solutions, including:

- pairs of local Dell controllers acting in an active-active mode or a hot-standby mode
- a master controller backing up a set of local controllers

VRRP eliminates a single point of failure by providing a mechanism among the controllers to elect a VRRP "master" ”controller. If VRRP preemption is disabled (the default setting) and all controllers share the same

priority, the first controller that comes up becomes the master. However, if you enable VRRP preemption and all controllers share the same priority, the controller with the highest IP address becomes the master.

The Virtual Router Redundancy Protocol (VRRP) is used to create various redundancy solutions, including pairs of local controllers acting in an active-active mode or a hot-standby mode, or a master controller backing up a set of local controllers. The master controller owns the configured virtual IP address for the VRRP instance.

When the master controller becomes unavailable, a backup controller steps in as the master and takes ownership of the virtual IP address. All network elements (APs and other controllers) can be configured to access the virtual IP address, thereby providing a transparent redundant solution to your network.

For more information on configuring the VRRP-Based Redundancy, refer to Configuring VRRP Redundancy on page 579.

High Availability Deployment Models

High availability supports the following deployment modes.

- Active/Active Deployment Model on page 571
- 1:1 Active/Standby Deployment Model on page 571
- N:1 Active/Standby Deployment Model
- Master-Redundancy Deployment Model

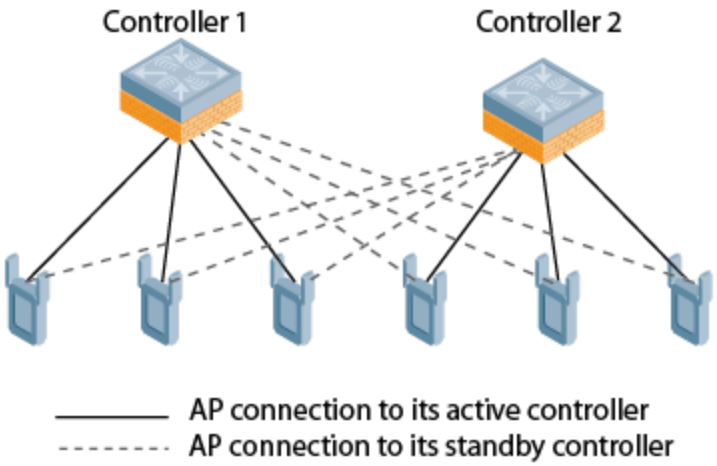


The High Availability Fast Failover feature supports APs in campus mode using tunnel, decrypt-tunnel, or bridge forwarding modes. This feature is not supported on remote APs and mesh APs in any mode.

Active/Active Deployment Model

In this model, two controllers are deployed in dual mode. Controller one acts as standby for the APs served by controller two, and vice-versa. Each controller in this deployment model supports approximately 50% of its total AP capacity, so if one controller fails, all the APs served by that controller would fail over to the other controller , thereby providing high availability redundancy to all APs in the cluster.

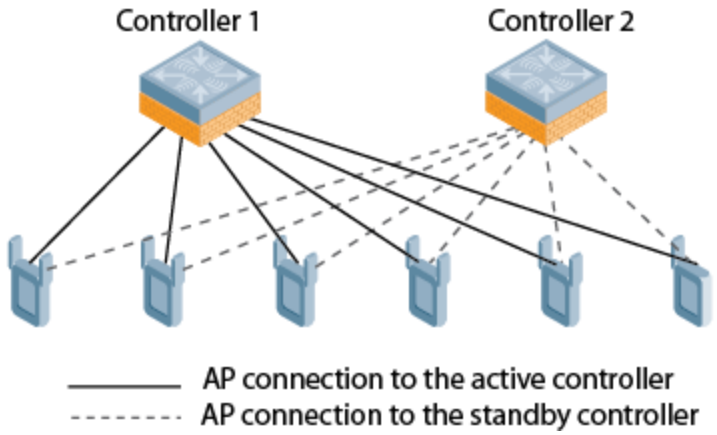
Figure 58 *Active-Active HA Deployment*



1:1 Active/Standby Deployment Model

In this model, the active controller supports up to 100% of its rated capacity of APs, while the other controller in standby mode is idle. If the active controller fails, all APs served by the active controller would failover to the standby controller.

Figure 59 *1:1 Active/Standby Deployment*

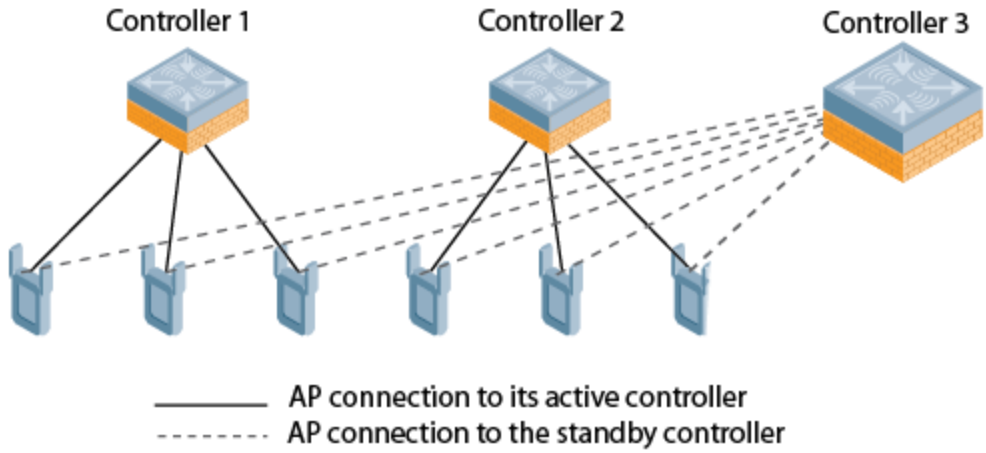


N:1 Active/Standby Deployment Model

In this model, the active controller supports up to 100% of its rated AP capacity, while the other controller is idle in standby mode. If an active controller fails, all APs served by the active controller would failover to the standby controller. This model requires that the AP capacity of the standby controller is able to support the total number of APs distributed across all active controllers in the cluster.

In the cluster shown in the example below, the standby controller has enough AP capacity to support the total number of APs terminating at the active controllers (Controller 1 and Controller 2).

Figure 60 *1:1 Active/Standby Deployment*

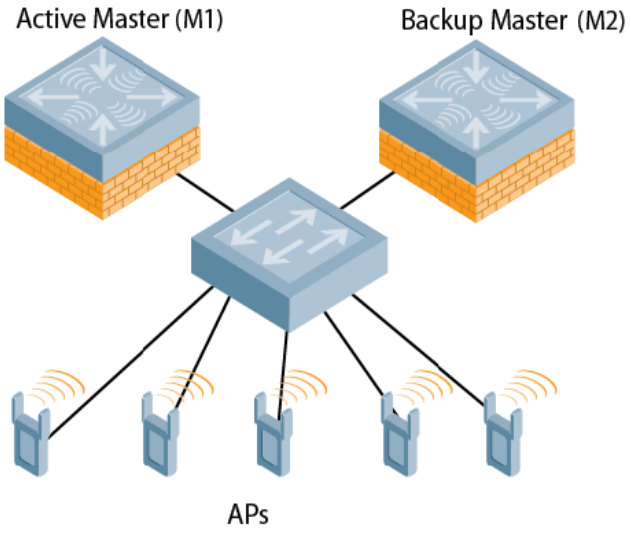


Master-Redundancy Deployment Model

ArubaOS supports VRRP-based LMS redundancy in a deployment with master-master redundancy. In the topology below, when an AP connects to the master controller (M1), the AP receives a standby IP, which it uses to establish a standby connection to the backup master (M2). If the active master becomes unreachable or reboots, the backup master changes its VRRP role to master, and accepts AP active connections.

When M1 comes back up, it initially acts as a backup master, and APs associated to M2 establish a standby connection to M1. When the controllers change roles and M1 becomes the active master once again, M2 forces APs to use M1 as their active master. If an AP has not established a connection to M1 before it disassociates from M2, the AP rebootstraps before it reconnects back to M1.

Figure 61 *Redundancy with a Active-Backup Master Controller Pair*



When a VRRP instance is configured on the controller vlan, there would be no change in the VRRP state if the failover scenario is tested by shutting down the port or bringing down the vlan. The controller continues to remain in the Master state and sends VRRP advertisements, which do not reach the peer controller. When the port is down the peer controller becomes the Master, but when the port on the previous master is enabled it takes over the Master state. The peer controller moves out of the master state when the original master sends a higher priority advertisement even when preemption is not enabled. The peer controller will not be preempted if the master controller crashes or reboots.

AP Communication with Controllers

The High Availability features work across Layer-3 networks, so there is no need for a direct Layer-2 connection between controllers in a high-availability group.

When the AP first connects to its active controller, the active controller provides the IP address of a standby controller, and the AP attempts to establish a tunnel to the standby to the standby controller. If an AP fails to connect to the first standby controller, the active controller will select a new standby controller for that AP, and the AP will attempt to connect to that standby controller.

An AP will failover to its backup controller if it fails to contact its active controller through regular heartbeats and keepalive messages, or if the user manually triggers a failover using the WebUI or CLI.

High Availability for bridge mode is supported on Campus APs. In this mode, the controller sends ACL Names to the APs instead of the ACL IDs. These APs generate and maintain the mapping between the ACL Name and ACL Id . In the event of a failover the ACL Name is sent to the AP from the stand-by controller. Since AP maintains the mapping, the ACL Ids remain intact during a failover.

Client State Synchronization

Client state synchronization allows faster client reauthentication in the event of a controller failure by synchronizing PMK and Key cache entries between active and standby controllers, When you enable this feature, clients only have to perform a four-way key exchange to reconnect to the network, (instead of performing a full authentication to the RADIUS server) dramatically shortening the time required for the client to reconnect.

The following section of this document describes topologies, guidelines, and limitations for this feature. **For the procedure to enable the client state synchronization feature, see** **Configuring High Availability on page 576****,**

Feature Guidelines and Limitations

Note the following guidelines and limitations before enabling this feature in your high availability deployment.

- W-7200 Series and W-3600 controllers, and W-6000M3 controller modules support this client state synchronization. This feature is not supported by W-3200, W-3400, and W-600 Series controllers.
- Only APs that support 802.11n and 802.11ac support client state synchronization.
- The client state synchronization and standby controller oversubscription features are mutually incompatible, and cannot be enabled simultaneously. If your deployment uses the standby controller oversubscription feature, you must disable it before you enable state synchronization.

High Availability Inter-Controller Heartbeats

The high availability inter-controller heartbeat feature allows for faster AP failover from an active controller to a standby controller, especially in situations where the active controller reboots or loses connectivity to the network.

The inter-controller heartbeat feature works independently from the AP mechanism that sends heartbeats from the AP to the controller. If enabled, the inter-controller heartbeat feature supersedes the AP's heartbeat to its controller. As a result, if a standby controller detects missed inter-controller heartbeats from the active controller, it triggers its standby APs to failover to the standby controller, *even if those APs have not detected any missed heartbeats between the APs and their active controller*. Use this feature with caution in deployments where the active and standby controllers are separated over high-latency WAN links.

When this feature is enabled, the standby controller starts sending regular heartbeats to an AP's active controller as soon as the AP has an UP status on the standby controller. The standby controller initially flags the active controller as **unreachable**, but changes its status to **reachable** as soon as the active controller sends a heartbeat response. If the active controller later becomes unreachable for the number of heartbeats defined by the heartbeat threshold (by default, 5 missed heartbeats), the standby controller immediately detects this error, and informs the APs using the standby controller to fail over from the active controller to the standby controller. If, however, the standby controller never receives an initial heartbeat response from the active controller, and therefore never marks the active controller as initially reachable, the standby controller will not initiate a failover.

This feature is disabled by default. It can be used in conjunction with the high availability state synchronization feature only in topologies that use a single active and standby controller, or a pair dual-mode active controllers that act as standby controllers for each other. High availability inter-controller heartbeats can be enabled and configured in the high-availability group profile using the WebUI or Command-Line interfaces.

For the procedure to enable and configure inter-controller heartbeats, see Configuring High Availability on page 576.

High Availability Extended Controller Capacity

The standby controller over-subscription feature allows a standby controller to support connections to standby APs beyond the controller's original rated AP capacity. This feature is an enhancement from the high availability feature introduced in ArubaOS 6.3.0.0, which requires the standby controller have a AP capacity equal to or greater than the total AP capacity of all the active controllers it supports.



The following section of this document gives an lists requirements and capacity limitations for this feature. For the procedure to enable the extended standby controller capacity, see Configuring High Availability on page 576,

The following section of this document gives an lists requirements and capacity limitations for this feature. For the procedure to enable the extended standby controller capacity, see Configuring High Availability on page 576,

Starting with ArubaOS 6.4.0.0, a W-7200 Series controller acting as a standby controller can oversubscribe to standby APs by up to four times that controller's rated AP capacity, and a standby W-6000M3 controller module or W-3600 controller can oversubscribe by up to two times its rated AP capacity, as long as the tunnels consumed the standby APs do not exceed the maximum tunnel capacity for that standby controller.

Feature Requirements

This feature can be enabled on controllers in a master-local topology where centralized licensing is enabled on the active and standby controllers, or on independent master controllers that are not using VRRP-based redundancy. If centralized licensing is disabled, the standby AP oversubscription feature are disabled also. Standby controller oversubscription and the high availability state synchronization features are mutually incompatible cannot be be enabled simultaneously. If your deployment uses the state synchronization feature, you must disable it before you enable standby controller oversubscription.



W-3200, W-3400 and W-600 Series controllers do not support this feature.

Standby Controller Capacity

The following table describes the AP oversubscription capacity maximum supported tunnels and for controllers that support this feature.

Table 105: *Controller Support for Standby Oversubscription*

Controller Model	Standby AP Capacity	Maximum Tunnels Supported
W-6000M3	2x rated AP capacity	16384 tunnels
W-3600	2x rated AP capacity	16384 tunnels
W-7210	4x rated AP capacity	16384 tunnels
W-7220	4x rated AP capacity	32768 tunnels
W-7240	4x rated AP capacity	65536 tunnels

To determine the number of standby tunnels consumed by APs on each active controller, multiply the number of APs on the active controllers by the number of BSSIDs per AP. As an example, consider a deployment with four active W-7210 controllers that each have 512 APs with 8 BSSIDs. The APs on each active controller consume (512 * 8) tunnels, for a combined total of 16,384 tunnels. A single W-7210 controller using the standby controller oversubscription feature can act as the standby controller for all four active controllers in this example, because this topology is within the 4x rated AP capacity limit and maximum tunnel limit for the W-7210 controller model.

If the network administrator later changed all the APs in this deployment to support 10 BSSIDs, each active controller would use (512 * 10) tunnels, for a combined total of 20,480 tunnels on the four active controllers. The tunnels required by the APs on the active controllers would then exceed the maximum tunnel limit for the standby controller, so the standby controller can no longer support all APs on the active controllers. Dynamic changes to configuration (such as the addition of BSSIDs to any AP group) causes all the standby APs to disconnect and reconnect back to the standby controller defined by their updated configuration

To view information about the numbers of currently associated APs and supported BSS tunnels, and the remaining capacity for additional APs and BSS tunnels, issue the CLI command **show ha oversubscription statistics.**

AP Failover

If a standby controller reaches its AP oversubscription capacity or exceeds its maximum BSSID limit, the standby controller drops any subsequent standby AP connections. A dropped AP attempts to reconnect to the standby controller, but after it exceeds the maximum number of request retries, the AP informs the active controller that it is unable to connect to the standby controller. The active controller then prompts the AP to create a standby tunnel to another standby controller, if one is configured.

If an active controller fails, the APs on the active controller fail over to the standby controller. Once the standby controller has reached its capacity for active APs,it terminates tunnels to any standby APs that controller can no longer serve. When these APs detect that there is no longer a heartbeat between the AP and the standby controller, they notify their active controller that they can no longer connect to the standby. The active controller then prompts the APs to establish standby tunnels to another standby controller, if one is configured.

Configuring High Availability

A controller using this feature can have one of three high availability roles: **active**, **standby** or **dual**. An active controller serves APs, but cannot act as a failover standby controller for any AP except for those that it serves as active. A standby controller acts as a failover backup controller, but cannot be configured as the primary controller for any AP. A dual controller can support both roles, acting as the active controller for one set of APs, and as a standby controller for another set of APs.



Starting with ArubaOS 6.4, a controller is assigned the **dual** role if no other role is specified

The high availability feature supports redundancy models with an active controller pair, or an active/standby deployment model with one backup controller supporting one or more active controllers. Each of these clusters of active and backup controllers comprises a high-availability group. All active and standby controllers within a single high-availability group must be deployed in master-local or independent masters topology. If you use an independent masters topology requires all independent master controllers to have the same WLAN configuration.

Pre-Deployment Information

Refer to the following sections of this documet for deployment models and feature details to help you plan your high availability solution,:

- High Availability Deployment Models on page 571
- High Availability Extended Controller Capacity on page 574
- Client State Synchronization on page 573
- High Availability Inter-Controller Heartbeats on page 574

Configuring High Availability

Configure the high availability feature in the WebUI or CLI using the high-availability and high-availability group profiles.

In the WebUI

To configure High Availability using the WebUI:

1. Navigate to **Configuration > Advanced Services > Redundancy**.
2. In the HA group information section in the right window pane, click **Add New**. A popup window appears
3. In the **Name** field, enter a name for the HA group you just created.
4. In the controller IP address field, enter the IP address of a controller in the HA group.
5. Click the IP Version drop-down list and select either IP4 or IPv6 to identify the IP address version type used by the controller.
6. Click the **Role** drop-down list, and assign a role to the controller. The IP address of each controller must be reachable by APs, and must be the IP address that appears in the **Configuration > Controller > System settings** tab of the controller WebUI, or in the output of the show controller-ip CLI command.
   - **Active**: Controller is active and is serving APs.
   - **Dual**: Controller serves some APs and acts as a standby controller for other APs.
   - **Standby**: Controller does not serve APs, as only acts as a standby in case of failover.
7. Click **Add** to add the controller to the group.
8. (Optional) The high availability inter-controller heartbeat feature allows for faster AP failover from an active controller to a standby controller, especially in situations where the active controller reboots or loses connectivity to the network.To edit the default heartbeat threshold and interval values:
   a. Enter a heartbeat threshold in the **Heartbeat Threshold** field to define the number of heartbeats that must be missed before the APs are forced to fail over to the standby controller. This value must be between 3 and 10, inclusive.
   b. Enter a heartbeat interval in the **Heartbeat Interval** field to define how often inter-controller heartbeats are sent. This value must be between 100 and 1000 ms, inclusive.
9. (Optional) State synchronization improves failover performance by synchronizing client authentication state information from the active controller to the standby controller. To use the state synchronization feature, enter a pre-shared key into the **Pre-shared key** field. Note, however, that this feature will not be enabled until you complete the task in step 13 on page 577
10. Click **OK**. The popup window closes, and the name of the new HA group appears in the **HA Group Configuration** field on the **Configuration > Advanced Services > Redundancy** page.
11. (Optional) Select the **Preemption** checkbox to require AP that has failed over to a standby to attempt to connect back to its original active controller once that controller is reachable again. When you enable this setting, the AP will wait for the time specified by the **lms-hold-down-period** parameter in the **ap system** profile before the AP attempts to switch back from the standby controller to the orginal controller.
12. (Optional) The standby controller oversubscription feature allows a standby controller to support connections to standby APs beyond the controller's original rated AP capacity. To enable this feature, click the **Oversubscription** checkbox.
13. (Optional) if you defined a pre-shared key in step 9 on page 577, select the **State Synchronization** checkbox to enable this feature. (For more information about State Synchronization, see Client State Synchronization on page 573.)
14. (Optional) The inter-controller heartbeat feature allows for faster AP failover from an active controller to a standby controller by enabling regular heartbeats between a standby controller and an active controller
15. Click **Apply**.

In the CLI

To configure a High Availability group using the command-line interface, access the CLI in config mode and issue the following commands. The high availability group profile should be configured with a pair of IPv4 controller addresses and pair of IPv6 controller addresses to allow an IPv4 or IPv6 access point to establish a connection to a standby controller.

```
ha group-profile <profile>
```

```
clone <profile-name>
controller <ipv4-ip-addr> role active|dual|standby
controller-v6 <ipv6-ip-addr> role active|dual|standby
heartbeat
heartbeat-interval <heartbeat-interval>
heartbeat-threshold <heartbeat-threshold>
no ...
over-subscription
pre-shared-key <key>
preemption
state-sync
```

A controller using the high availability features must be defined as a member of a high availability group.To add a controller to the new high availability group, issue following CLI command:

```
(host)(config)#ha group-membership <ha-group>
```

Migrating from VRRP or Backup-LMS Redundancy

ArubaOS has a concept of a local management switch (LMS) and a backup LMS. In a typical deployment, the AP contacts the master mobility controller and is directed to the mobility controller that handles the AP connection and traffic via the **LMS** parameter. If the LMS becomes unreachable and a backup LMS is specified, the AP attempts to reconnect to that backup mobility controller. This function provides Layer 3 and site redundancy when this level of redundancy is required.



High Availability: Fast Failover provides redundancy for APs, but not for controllers. Deployments that require master controller redundancy should continue to use an existing VRRP redundancy solution.

If your deployment currently uses a backup-LMS or VRRP redundancy solution, use the procedures below to migrate to a High-Availability based solution. For more information on this topology, see also Master-Redundancy Deployment Model on page 572

Configuring a Master Controller for Redundancy and High Availability:

Starting with ArubaOS 6.4, a backup-master controller can use the High Availability feature. However a backup-master controller can only accept standby connections from APs, and will not serve active APs as long as its master redundancy role is **backup**.

This type of High Availability deployment has the following requirements and limitations :

- A backup-master controller can only form an active-standby pair with the master controller.
- The backup master cannot terminate active APs.
- Both the backup-master and master controllers must be configured with the **dual** controller role.
- The controller IP address defined in the high availability group profile must be the IP address of the VRRP interface.
- **The inter-controller heartbeat feature is not recommended for backup-master and master controller pairs using the High Availability feature.** If the inter-controller heartbeat feature is enabled in a high availability group profile for redundant masters, the inter-controller failover time must be greater than the VRRP failover time. That is, the (heartbeat interval * heartbeat threshold threshold) value should be greater than the (advertisement time * 3 + preemption delay + skew time [which is based on priority]).

Perform the following steps to configure controller high availability on a backup-master and master controller pair.

1. Configure the high-availability group profile with a **dual** role for the master controller.

```
(host)(config) #ha group-profile grp1
```

```
(host)(HA group information "grp1"): controller <VRRP interface ipaddress> role dual
```

2. Configure the high-availability group profile with a **dual** role for the backup-master controller.

```
(host) (HA group information "grp1"): controller <VRRP interface ipaddress> role dual
```

Migrating from VRRP Redundancy

Perform the following steps to migrate from VRRP to High-Availability redundancy:

1. Remove the VRRP IP address as the LMS IP address of the AP.

```
(host) (AP system profile) #no lms-ip
```

2. Configure the AP to use the active controller's IP address (not VRRP the IP address) as the LMS-IP for the AP.

```
(host) (AP system profile) #lms-ip <ipaddress>
```

3. Configure the AP to use the standby controller IP address (not VRRP the IP address) as the backup LMS-IP for the AP.

```
(host) (AP system profile) #bkup-lms-ip <ipaddress>
```

4. Configure the master controller with a dual role in the high-availability group profile.

```
(host) (config) #ha group-profile grp1
(host) (HA group information "grp1"): controller <ipaddress> role dual
```

5. Configure the standby controller with a dual role in the high-availability group profile.

```
(host) (HA group information "grp1"): controller <ipaddress> role dual
```

Migrating from Backup-LMS Redundancy

Perform the following steps to migrate from Backup-LMS to High-Availability redundancy and maintain the existing configuration as defined by the **lms-ip** and **bkup-lms-ip** parameters in the AP system profile.

1. Configure the controller serving the AP with a dual role in the high-availability group profile.

```
(host) (config) #ha group-profile grp1
(host) (HA group information "grp1"): controller <ipaddress> role dual
```

2. Configure the AP's standby controller with a dual role in the high-availability group profile.

```
(host) (HA group information "grp1"): controller <ipaddress> role dual
```

Configuring VRRP Redundancy

In a Dell network, the APs are controlled by a controller. The APs tunnel all data to the controller which processes the data, including encryption/decryption and bridging/forwarding. Local controller redundancy refers to providing redundancy for a controller such that the APs fail over"to a backup controller if a controller becomes unavailable. Local controller redundancy is provided by running VRRP between a pair of controllers.The APs are then configured to connect to the "virtual-IP" configured for the VRRP instance.



The two controllers need to be connected on the same broadcast domain (or Layer-2 connected) for VRRP operation. The two controllers should be of the same class (for example, and both should be running the same version of ArubaOS.

The following section of this document includes the following procedures:

- Before you Begin on page 580
- Configuring the Local Controller for Redundancy on page 580
- Configuring the LMS IP on page 582
- Configuring the Master Controller for Redundancy on page 582
- Configuring Database Synchronization on page 584
- Enabling Incremental Configuration Synchronization (CLI Only) on page 584
- Configuring Master-Local Controller Redundancy on page 585

Before you Begin

Before you start the configuring VRRP redundancy obtain the following network information:

- **VLAN ID** on the two local controllers that are on the same Layer-2 network.
- **Virtual IP address** to be used for the VRRP instance.

Configuring the Local Controller for Redundancy

You can use either the WebUI or CLI to configure VRRP on the local controllers. For this topology, it is recommended to use the default priority value.

In the WebUI

1. Navigate to the **Configuration > Advanced Services > Redundancy** page on the WebUI for each of the local controllers.
2. Under Virtual Router Table, click **Add** to create a VRRP instance.
3. Select the IP version.
4. Enter the IPv4\IPv6 Address for the virtual router. Select the VLAN on which VRRP will run. Set the **Admin State** to **Up**.
5. Configure other VRRP parameters as described in the table below.
6. Click **Done**.

Table 106: *VRRP Parameters*

Parameter	Description
IP Version	Select IPv4 \ IPv6 from the drop-down list box.
Virtual Router ID	This uniquely identifies this VRRP instance. For ease in administration, you should configure this with the same value as the VLAN ID.
Advertisement Interval (secs)	This is the interval, in seconds, between successive VRRP advertisements sent by the current *master*. The default interval time is recommended. Default: 1 second
Authentication Password	This is an optional password of up to eight characters that can authenticate VRRP peers in their advertisements. If this is not configured, there is no authentication password set.
Description	This is an optional text description to describe the VRRP instance.
IP \ IPv6 Address	Based on the selection made in the IP version field, either IP Address \ IPv6 Address is displayed. This is the virtual IP address that will be owned by the elected VRRP master. Ensure that the same IP address and VRRP ID is used on each member of the redundant pair. Note: The IP address must be unique and cannot be the loopback address of the controller. A maximum of only two virtual IPv6 addresses can be configured on each VRRP instance. Only IPv6 address format is supported for the v6 instance.
Enable Router Pre-emption	Selecting this option means that a controller can take over the role of *master* if it detects a lower priority controller currently acting as *master.*
Delay	Specifying a value enables the delay timer. The timer is triggered when the VRRP state moves out of backup or init state to become a master. This is applicable only if you enable router pre-emption.

Table 106: *VRRP Parameters*

Parameter	Description
	When the timer is triggered, it forces VRRP to wait for a specified period of time, so that all the applications are ready before coming up. This prevents the APs from connecting to the controller before it can receive them. In the mean time, if there is an advertisement from another VRRP, the VRRP stops the timer and does not transition to master.
Priority	Priority level of the VRRP instance for the controller. This value is used in the election mechanism for the *master.*
Tracking	Configures a tracking mechanism that modifies a specified *value* to the priority after a controller has been the master for the VRRP instance. This mechanism is used to avoid failing over to a backup Master for transient failures. Tracking can be based on one of the following: • Master Up Time: how long the controller has been the master. The value of *duration* is the length of time that the administrator expects will be long enough that the database gathered in the time is too important to be lost. This will obviously vary from instance to instance. • VRRP Master State Priority: the master state of another VRRP. Tracking can also be based on the interface states of the controller: • VLAN and Interface: prevents asymmetric routing by tracking multiple VRRP instances. The priority of the VRRP interface determined by the *sub* value can increase or decrease based on the operational and transitional states of the specified VLAN or Fast Ethernet/Gigabit Ethernet port. When the VLAN or interface comes up again, the value is restored to the previous priority level. You can track a combined maximum of 16 interfaces and VLANs. For example, you can track an interface that connects to a default gateway. In this situation, configure the VRRP priority to decrease and trigger a VRRP master re-election if the interface goes down. This not only prevents network traffic from being forwarded, but reduces VRRP processing.
Admin State	Administrative state of the VRRP instance. To start the VRRP instance, change the admin state to **UP** in the WebUI.
VLAN	VLAN on which the VRRP protocol runs.

In the CLI

```
(host)(config)#vrrp <id>
   advertise <interval>
   authentication <password>
   description <text>
   ip address <ipaddr>
   no...
   preempt
   priority <level>
   shutdown
   tracking interface {fastethernet <slot>/<port>|gigabitethernet <slot>/<port>}
     {sub <value>}
   tracking master-up-time <duration> add <value>
   tracking vlan <vlanid> {sub <value>}
   tracking vrrp-master-state <vrid> add <value>
   vlan <vlanid>

(host) (config)#vrrp ipv6 <id>
   advertise <interval>
   description <text>
   ipv6 address <ipaddr>
   no...
   preempt
```

```
priority <level>
shutdown
tracking interface {fastethernet <slot>/<port>|gigabitethernet <slot>/<port>}
  {sub <value>}
tracking master-up-time <duration> add <value>
tracking vlan <vlanid> {sub <value>}
tracking vrrp-master-state <vrid> add <value>
vlan <vlanid>
```

Configuring the LMS IP

Configure the APs to terminate their tunnels on the virtual-IP address. To specify the controller to which an AP or AP group tunnels client traffic, you configure the LMS IP in the AP system profile on the master controller. For information on how to configure the LMS IP in the AP system profile, see Optional AP Configuration Settings on page 503



This configuration needs to be executed on the master controller; the APs obtain their configuration from the master controller.

In the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration** page on the master controller.
   - If you select **AP Group**, click **Edit** for the AP group name for which you want to configure the LMS IP.
   - If you select **AP Specific**, select the name of the AP for which you want to configure the LMS IP.
2. Under the **Profiles** section, select **AP** to display the AP profiles.
3. Select the AP system profile you want to modify.
4. Enter the controller IP address in the LMS IP field.
5. Click **Apply**.

In the CLI

On the master controller:

```
(host)(config) #ap system-profile <profile>
  lms-ip <ipaddr>

(host)(config) #ap-group <group>
  ap-system-profile <profile>

(host)(config) #ap-name <name>
  ap-system-profile <profile>
```

Configuring the Master Controller for Redundancy

The master controller in the Dell user-centric network acts as a single point of configuration for global policies such as firewall policies, authentication parameters, and RF configuration to ease the configuration and maintenance of a wireless network. It also maintains a database related to the wireless network that you can use to make adjustments (automated as well as manual) in reaction to events that cause a change in the environment (such as an AP becoming unavailable).

The master controller is also responsible for providing the configuration for any AP to complete its boot process. If the master controller becomes unavailable, the network continues to run without any interruption. However, any change in the network topology or configuration will require the availability of the master controller.

To maintain a highly redundant network, the administrator can use a controller to act as a hot standby for the master controller. The underlying protocol used is the same as in local redundancy, that is, VRRP.

- Collect the following data before configuring master controller redundancy.
  - **VLAN ID** on the two controllers that are on the same layer 2 network.
  - **Virtual IP address** that has been reserved to be used for the VRRP instance.
- You can use either the WebUI or CLI to configure VRRP on the master controllers (see Table 106). For this topology, the following are recommended values:
  - For priority: Set the master to 110; set the backup to 100 (the default value)
  - Enable preemption
  - Configure master up time or master state tracking with an add value of 20.

The following is a configuration example for the initially-preferred master.

```
(host)(config) #vrrp 22
   vlan 22
   ip address 10.200.22.254
   priority 110
   preempt
   authentication password
   description Preferred-Master
   tracking master-up-time 30 add 20
   no shutdown
```

The following configuration is the corresponding VRRP configuration for the peer controller.

```
(host)(config) #vrrp 22
   vlan 22
   ip address 10.200.22.254
   priority 100
   preempt
   authentication password
   description Backup-Master
   tracking master-up-time 30 add 20
   no shutdown
```

Use the following commands to associate the VRRP instance with master controller redundancy.

Table 107: *VRRP Commands*

Command	Explanation
`master-redundancy`	Enter the master-redundancy context.
`master-vrrp <id>`	Associates a VRRP instance with master redundancy. Enter the virtual router ID of the VRRP instance.
`peer-ip-address <ipaddr> ipsec <key>`	Loopback IP address of the peer controller for master redundancy. The pre-shared key secures communication between the master controllers. Specify a key of up to 64 characters.
`masterip <ipaddr> ipsec <key>`	Configures the master IP address and pre-shared key on a local controller for communication with the master controller. Configure this to be the virtual IP address of the VRRP instance used for master redundancy.



Configure all the APs and local controllers in the network with the virtual IP address as the master IP address. You can configure the master IP address for local controllers during the Initial Setup. The controller will require a reboot after changing the master IP on the controller.

If DNS resolution is the chosen mechanism for the APs to discover their master controller, ensure that the name *"aruba-master"* resolves to the same virtual IP address configured as a part of the master redundancy.

Configuring Database Synchronization

In a redundant master controller scenario, you can configure a redundant pair to synchronize their WMS and local user databases. You can either manually or automatically synchronize the databases.

When manually synchronizing the database, the active VRRP master synchronizes its database with the standby. The command takes effect immediately.

When configuring automatic synchronization, you set how often the two controllers synchronize their databases. To ensure successful synchronization of database events, you should set periodic synchronization to a minimum period of 20 minutes.

In the WebUI

1. On each controller, navigate to the **Configuration > Advanced Services > Redundancy** page.
2. Under Database Synchronization Parameters, do the following:
   a. Select the **Enable periodic database synchronization** check box. This enables database synchronization.
   b. Enter the frequency of synchronizing the databases. A minimum value of 20 minutes is recommended.
3. Click **Apply**.

In the CLI

Use the following commands to configure database synchronization.

Table 108: *Database synchronization commands*

Command	Description
`database synchronize`	This enable mode command manually synchronizes the databases and takes effect immediately.
`database synchronize period <minutes>`	This config mode command defines the scheduled interval for synchronizing the databases.

To view the database synchronization settings on the controller, use the following command:

```
(host)#show database synchronize
```

Enabling Incremental Configuration Synchronization (CLI Only)

Typically when the master and the local is synchronized, the complete configuration is sent to the local. You can, now send only the incremental updates to the local using the following CLI commands.

Table 109: *Incremental Configuration Synchronization Commands*

Command	Description
`cfgm set sync-type <complete>`	The master sends full configuration file to the local.
`cfgm set sync-type <snapshot>`	The master sends only the incremental configuration to the local. **NOTE:** This configuration is enabled by default
`cfgm set sync-command-block <number>`	To configure the number of command-list blocks. Each block contains a list of global configuration commands for each write-mem operationThe number is 3 by default
`show master-configpending`	To show a list of global commands, which are not saved but are sent to the local.
`clear master-local-session <A.B.C.D>`	To manually push the full configuration to the local.

Configuring Master-Local Controller Redundancy

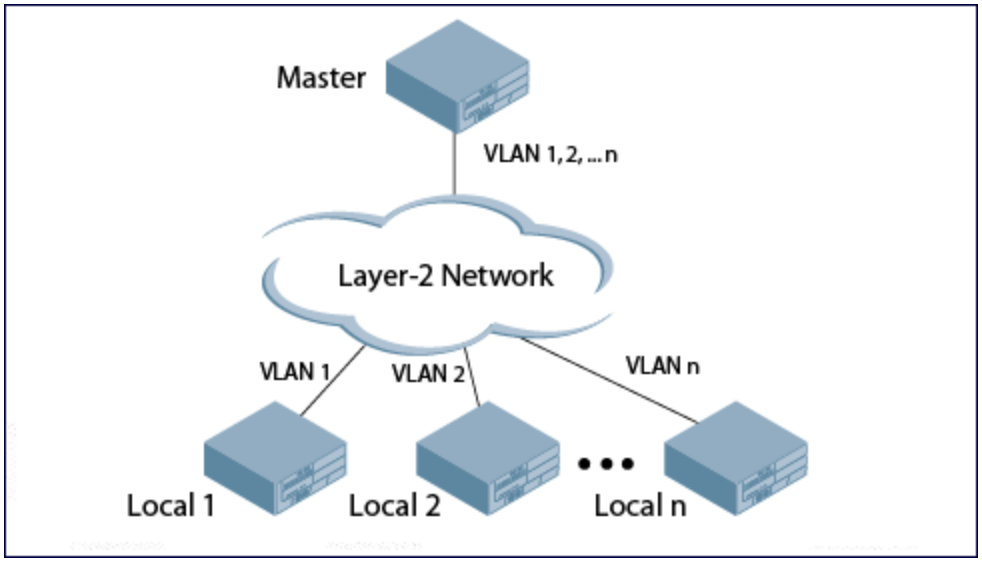
This section outlines the concepts behind a redundancy solution where a master can act as a backup for one or more local controllers, and shows how to configure the Dell controllers for such a redundant solution. In this solution, the local controllers act as the controller for the APs. When any one of the local controllers becomes unavailable, the master takes over the APs controlled by that local controller for the time that the local controller remains unavailable. It is configured such that when the local controller comes back again, it can take control over the APs once more.

This type of redundant solution is illustrated by the following topology diagram.



This solution requires that the master controller have Layer-2 connectivity to all the local controllers.

Figure 62 *Redundant Topology: Master-Local Redundancy*



The network in Figure 62, the master controller is connected to the local controllers on VLANs 1 through *n* through a Layer-2 network. To configure redundancy as described in the conceptual overview for master-local redundancy, configure VRRP instances on each of the VLANs between the master and the respective local

controller. The VRRP instance on the local controller is configured with a higher priority to ensure that when available, the APs always choose the local controller to terminate their tunnels.

- Configure the interface on the master controller to be a trunk port with 1, 2... *n* being member VLANs.
- Collect the following data before configuring master controller redundancy.
  - **VLAN IDs** on the controllers corresponding to the VLANs 1, 2...*n* shown in the topology above.
  - **Virtual IP addresses** that has been reserved to be used for the VRRP instances.
- You can use either the WebUI or CLI to configure VRRP on the master controllers (see Table 106). For this topology, the following are recommended values:
  - For priority: Set the local to 110; set the master to 100 (the default value)
  - Enable preemption



The master controller is configured for a number of VRRP instances (equal to the number of local controllers the master is backing up).

To configure APs, configure the appropriate virtual IP address (depending on which controller is expected to control the APs) for the LMS IP address parameter in the AP system profile for an AP group or specified AP.



Configure these AP settings on the master controller, not the local controller.

As an example, the administrator configures APs in the AP group “floor1” to be controlled by local controller 1, APs in the AP group “floor2” to be controlled by local controller 2 and so on. All the local controllers are backed up by the master controller. In the AP system profile for the AP group “floor1”, enter the virtual IP address (10.200.22.154 in the example configuration) for the LMS IP address on the master controller.

Configuration changes take effect only after you reboot the affected APs; this allows them to reassociate with the local controller. After rebooting, these APs appear to the new local controller as local APs.

The ArubaOS implementation of Rapid Spanning Tree Protocol (RSTP) is as specified in 802.1w, with backward compatibility to legacy Spanning Tree (STP) 802.1D. RSTP takes advantage of point-to-point links and provides rapid convergence of the spanning tree. RSTP is enabled by default on all Dell controllers.

Topics in this chapter include:

- Understanding RSTP Migration and Interoperability on page 587
- Working with Rapid Convergence on page 587
- Configuring RSTP on page 589
- Troubleshooting RSTP on page 591

Understanding RSTP Migration and Interoperability

The ArubaOS RSTP implementation interoperates with PVST (Per VLAN Spanning Tree 802.1D) and Rapid-PVST (802.1w) implementation on industry-standard router/switches. Dell supports global instances of STP and RSTP only. Therefore, the ports on industry-standard routers/switches must be on the default or untagged VLAN for interoperability with Dell controllers.

ArubaOS supports RSTP on the following interfaces:

- FastEthernet IEEE 802.3: fastethernet
- Gigabitethernet IEEE 802.3: gigabitethernet
- Port Channel ID: port-channel

Working with Rapid Convergence

Since RSTP is backward compatible with STP, it is possible to configure both the bridges in the same network. However, such mixed networks may not always provide rapid convergence. RSTP provides rapid convergence when interfaces are configured as either:

- Edge ports: These are the interfaces/ports connected to hosts. These interfaces are immediately moved to the forwarding state. In this mode, an interface forwards frames by default until it receives a BPDU (Bridge Protocol Data Units) indicating that it should behave otherwise. It does not go through the Listening and Learning states.
- Point-to-Point links: These are the interfaces/ports connected directly to neighboring bridges over a point-to-point link. RSTP negotiates with the neighbor bridge for rapid convergence/transition only when the link is point-to-point.

Table 110 compares the port states between STP and RSTP.

Table 110: *Port State Comparison*

STP (802.1d) Port State	RSTP (802.1w) Port State
Disabled	Discarding
Blocking	Discarding
Listening	Discarding
Learning	Learning
Forwarding	Forwarding

In addition to port state changes, RSTP introduces port roles for all the interfaces (see Table 111).

Table 111: *Port Role Descriptions*

RSTP (802.1w) Port Role	Description
Root	The port that receives the best BPDU on a bridge.
Designated	The port can send the best BPDU on the segment to which it is connected.
Alternate	The port offers an alternate path, in the direction of root bridge, to that provided by bridge's root port.
Backup	The port acts as a backup for the path provided by a designated port in the direction of the spanning tree.

The `show spantree` command (configuration mode) output reveals the state and port role:

```
(host) (config) #show spantree

Designated Root MAC       00:0b:86:50:3c:20
Designated Root Priority  32768
Root Max Age 20 sec   Hello Time 2 sec   Forward Delay 15 sec

Bridge MAC                00:0b:86:50:3c:20
Bridge Priority           32768
Configured Max Age 20 sec   Hello Time 2 sec   Forward Delay 15 sec

Rapid Spanning-Tree port configuration
------------------------------
Port    State       Cost  Prio  PortFast  P-to-P  Role
----    -----       ----  ----  --------  ------  ----
FE 1/0  Discarding  0     128   Disable   Enable  Disabled
FE 1/1  Forwarding  0     128   Disable   Enable  Designated
FE 1/2  Forwarding  0     128   Disable   Enable  Root
FE 1/3  Discarding  0     128   Disable   Disable Disabled
FE 1/4  Discarding  0     128   Disable   Enable  Alternate
```

Also, the `show spanning-tree interface` command indicates the state and roles; see the partial output below:

```
(host) #show spanning-tree interface fastethernet 1/1

Interface FE 1/7 (port 8) in Spanning tree is FORWARDING
Port path cost 19, Port priority 128 Role DESIGNATED
```

Edge Port and Point-to-Point

At the interface level, the `portfast` command specifies an interface as an edge port and the `point-to-point` command specifies an interface as a point-to-point link. Since RSTP is enabled by default, all the interfaces are point-to-point links by default.

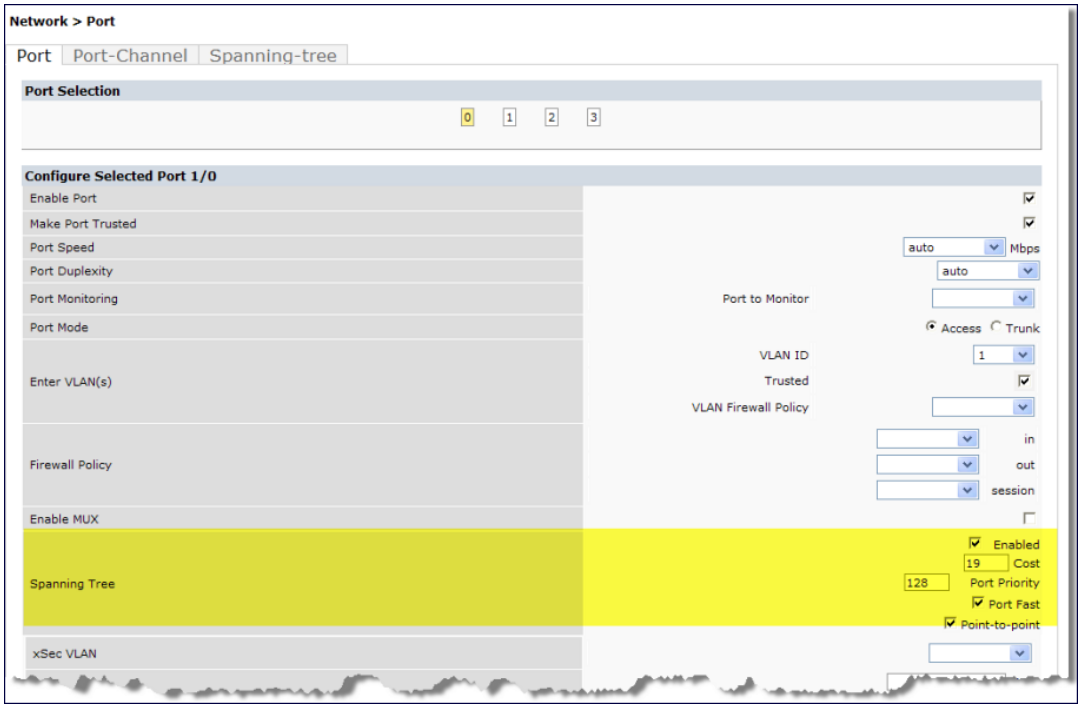
Configuring RSTP

Use either the command line reference or the WebUI to configure RSTP.

In the WebUI

The RSTP port interface is designated as point-to-point, by default, in the existing port configuration screen (Figure 63).

Figure 63 *Configuring RSTP*



Since RSTP is enabled by default, the default values appear in the WebUI. Table 112 lists the RSTP defaults and ranges (when applicable) in the configuration interface mode (config-if).

Table 112: *RSTP Default Values*

Feature	Default Value/Range
Port Cost	The RSTP interface path cost. Range: 1–65536 Default: Based on Interface type: Fast Ethernet 10Mbs: 100 Fast Ethernet 100Mbs: 19 1 Gigabit Ethernet: 4 10 Gigabit Ethernet: 2
Priority	Change the interface's RSTP priority Range: 0–255 Default: 128
Port Fast	Change from blocking to forwarding Default : Disabled
Point-to-Point	Set the interface as a point-to-point link Default : Enabled

In the CLI

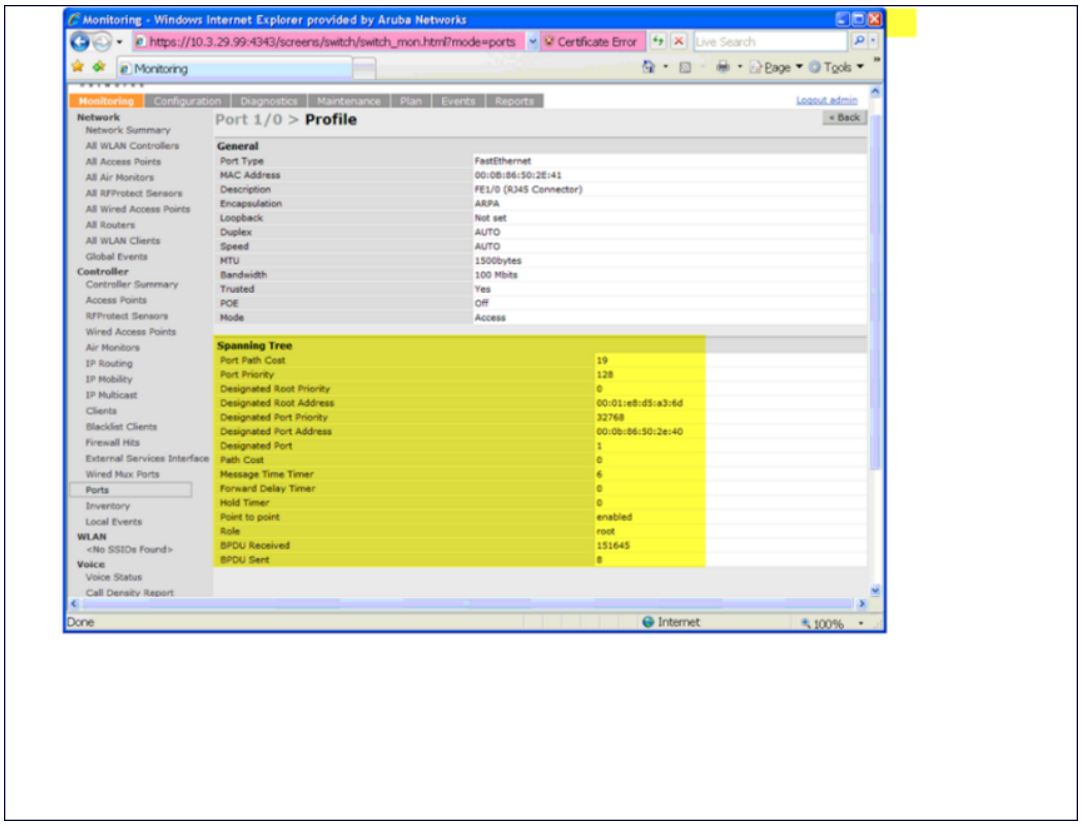
Change the default configurations via the command line.

```
(host) (config-if)#spanning-tree ?
cost                    Change an interface's spanning tree path cost
point-to-point          Set interface as point-to-point link
port-priority           Change an interface's spanning tree priority
portfast                Allow a change from blocking to forwarding
```

Monitoring RSTP

Statistical information for point-to-point, role, BPDU etc. can be viewed from the WebUI (see Figure 64).

Figure 64 *Monitoring RSTP*



Troubleshooting RSTP

The following points give some troubleshooting tips.

- The `show spantree` command displays the root and the bridge information; verify that they are correct. Also displayed is the port/interface information (for example, state, role, and so on); make sure that the state and role information correspond to each other.

```
(host) (config) #show spantree
Designated Root MAC       00:0b:86:50:3c:20
Designated Root Priority  32768
Root Max Age 20 sec   Hello Time 2 sec   Forward Delay 15 sec

Bridge MAC                00:0b:86:50:3c:20
Bridge Priority           32768
Configured Max Age 20 sec   Hello Time 2 sec   Forward Delay 15 sec

Rapid Spanning-Tree port configuration
------------------------------
Port    State       Cost  Prio  PortFast  P-to-P  Role
----    -----       ----  ----  --------  ------  ----
FE 1/0  Discarding  0     128   Disable   Enable  Disabled
FE 1/1  Forwarding  0     128   Disable   Enable  Designated
FE 1/2  Forwarding  0     128   Disable   Enable  Root
FE 1/3  Discarding  0     128   Disable   Disable Disabled
FE 1/4  Discarding  0     128   Disable   Enable  Alternate
```

- The `show spanning-tree interface` command (config-if mode) displays Tx/Rx BPDU counters. Validate those values. For example, if a port's role is "designated," it only transmit BPDUs and does not receive any.

In this case, the Tx counter will keep incrementing while Rx counter will remain the same. It is opposite for a port with role as “root/alternate/backup.”

```
(host) (config-if)#show spanning-tree interface fastethernet 1/1

Interface FE 1/1 (port 2) in Spanning tree is FORWARDING
Port path cost 19, Port priority 128 Role DISNIGNATED
PortFast DISABLED P-to-P ENABLED
Designated root has priority 0 address 00:01:e8:d5:a3:6d
Designated bridge has priority 32768 address 00:0b:86:50:58:30
Designated port is 2, path cost 0
Timers:  message age 0, forward delay 20, hold 0
Counts:  BPDUs received 0, sent 0

(host) (config-if)#
```

PVST+ (Per-VLAN Spanning Tree Plus) provides for load balancing of VLANs across multiple ports, resulting in optimal usage of network resources. PVST+ also ensures interoperability with industry accepted PVST+ protocols.

NOTE

PVST+ is disabled by default.

Topics in this chapter include:

- Understanding PVST+ Interoperability and Best Practices on page 593
- Enabling PVST+ in the CLI on page 593
- Enabling PVST+ in the WebUI on page 594

Understanding PVST+ Interoperability and Best Practices

The interoperability between RSTP and PVST+ are:

- When the access port on the controller and the trunk port terminate on one Layer 2 switch running PVST+, PVST+ will send untagged STP BPDUs on the access port; it also transmits untagged STP BPDUs (in addition to the other PVST+ BPDUs) on the native VLAN trunk port. If the Dell controller is the root, it will detect a loop on the native VLAN.

NOTE

If PVST+ is not on the controller, best practices recommends disabling RSTP on the Dell controller to avoid a looping issue.

- For VLAN load balancing when controllers are connected to armed mode, the VLAN priorities on two ports and bridge priorities must be configured so that one set of VLANs are active on one link, and the other set of VLANs are active on the other link.
- Supported instances are: 64 on the W-7200 Series, W-6000M3, and W-3000 Series, and 32 on the W-600 Series.

Enabling PVST+ in the CLI

PVST+ is disabled by default. Enable PVST+, ensure a VLAN instance is configured, and then configure PVST+.

1. Enable PVST+:
   ```
   spanning-tree mode rapid-pvst
   ```
2. Configure PVST+ forward time; the following command sets the time VLAN 2 spends in the listening and learning state (3 seconds):
   ```
   spanning-tree vlan 2 forward-time 3
   ```
3. Configure PVST+ hello time; the following command sets the time VLAN 2 waits to transmit BPDUs to four seconds:
   ```
   spanning-tree vlan 2 hello-time 4
   ```
4. Configure PVST+ max age; the following command sets the time VLAN 2 waits to receive a hello packet to 30 seconds:
   ```
   spanning-tree vlan 2 max-age 30
   ```

5. Configure PVST+ priority: the following command sets the VLAN 2 priority to 10, making it more likely to become the root bridge:

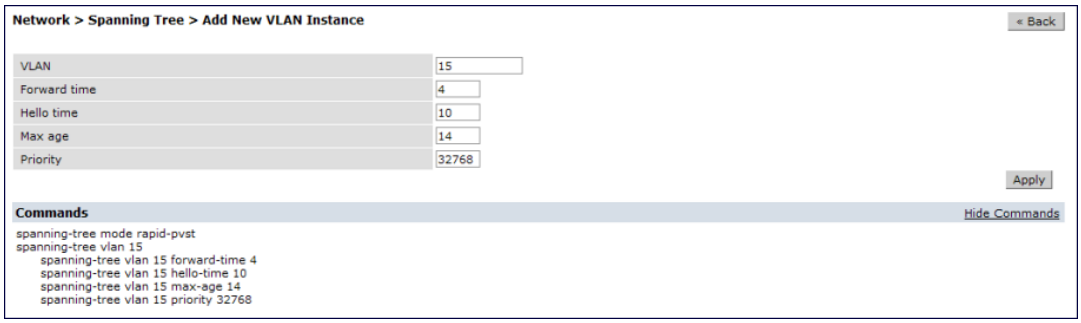
   ```
   spanning-tree vlan 2 priority 10
   ```

6. Configure PVST+ on a range of VLANs using the VLAN IDs (coma separated or hyphen separated):

   ```
   spanning-tree vlan range 2-6,11
   ```

Enabling PVST+ in the WebUI

From the WebUI, add a VLAN instance and enable PVST+:



ArubaOS provides support for Link Layer Discovery Protocol (LLDP) on the controllers to advertise identity information and capabilities to other nodes on the network, and store the information discovered about the neighbors.

This chapter contains the following major sections:

- Important Points to Remember on page 595
- LLDP Overview on page 595
- Configuring LLDP on page 596
- Monitoring LLDP Configuration on page 596

Important Points to Remember

- Inventory-management, and Location TLVs are not currently supported.
- Aggregation-management and Power-management TLVs are not supported.
- LLDP-MED will be supported in a future release.
- SNMP support is currently not available for LLDP MIBs.
- LLDP is not supported on the expanded slots and the management port of the W-6000M3 controller.
- Cisco Discovery Protocol (CDP) proprietary is not supported.
- The maximum number of neighbors that can be learnt on the controllers (including all the per port neighbor) is 250.

LLDP Overview

Link Layer Discovery Protocol (LLDP), defined in the IEEE 802.1AB standard, is a Layer 2 protocol that allows network devices to advertise their identity and capabilities on a LAN. ArubaOS supports a simple one-way neighbor discovery protocol with periodic transmissions of LLDP PDU.

- LLDP frames are constrained to a local link.
- LLDP frames are TLV (Type-Length-Value) form.
- LLDP Multicast address is 01-80-C2-00-00-0E.

LLDP provides support for a set of attributes used to discover neighbor devices. These attributes are referred as TLVs which contain type, length, and value descriptions. LLDP supported devices use TLVs to receive and send information such as configuration information, device capabilities, and device identity to their neighbors.

ArubaOS supports the following optional basic management TLVs which are enabled by default:

- MAC Phy configuration TLV
- Management address TLV
- Maximum frame size TLV
- Port-description TLV
- Port VLAN ID TLV
- System capabilities TLV
- System description TLV

- System name TLV
- VLAN name TLV

Default LLDP Configuration

To display the default LLDP information, use the following command:

```
(host) #show lldp interface gigabitethernet 1/1
Interface: FE1/1
LLDP Tx: Disabled, LLDP Rx: Disabled
Proprietary Neighbor Discovery: Disabled
LLDP-MED: Disabled
Fast Transmit interval: 1, Fast Transmit message counter: 4
Transmit interval: 30, Transmit hold 4, Hold timer: 120
```



When you use the default LLDP configuration, the **RX** and **TX** parameters are disabled. You have to explicitly enable them for LLDP to work.

Configuring LLDP

You can configure LLDP using the CLI. For detailed information on the LLDP commands, see the *ArubaOS Command Line Reference Guide*.

```
(host)(config) #interface gigabitethernet <slot/port>
(host)(config-if) #lldp
   fast-transmit-counter <1-8>
   fast-transmit-interval <1-3600>
   med
   receive
   transmit
   transmit-hold <1-100>
   transmit-interval <1-3600>
```



In the current release, the **med** command is not supported and will be supported in a future release. Also, the **fast-transmit-interval <1-3600>** and the **fast-transmit-counter <1-8>** parameters should be used with **med** command.

Monitoring LLDP Configuration

This section describes commands for monitoring LLDP configurations.

Display LLDP Interface

To display all LLDP information for all interfaces, use the following command:

```
(host)# show lldp interface
LLDP Interfaces Information
---------------------------
Interface  LLDP TX  LLDP RX  LLDP-MED  TX interval  Hold Timer
---------  -------  -------  --------  -----------  ----------
GE1/3      Enabled  Enabled  Enabled   30           120
GE1/4      Enabled  Enabled  Enabled   30           120
```

Display LLDP Interface <interface>

To display LLDP information for a specific interface, use the following command:

```
(host) #show lldp interface gigabitethernet <slot/port>
(host) #show lldp interface gigabitethernet <1/3>
Interface: gigabitethernet0/0/1
```

```
LLDP Tx: Enabled, LLDP Rx: Enabled
Proprietary Neighbor Discovery: Disabled
LLDP-MED: Disabled
Fast Transmit interval: 1, Fast Transmit message counter: 4
Transmit interval: 30, Hold timer: 120
```

Display LLDP Neighbor

```
(host)#show lldp neighbor
Capability codes: (R)Router, (B)Bridge, (A)Access Point, (P)Phone, (O)Other
LLDP Neighbor Information
-------------------------
Local Intf Chassis ID        Capability  Remote Intf  Expiry-Time (Secs)
---------  -----------       ----------  -----------  ------------------
GE1/3      00:0b:86:6a:25:40  B:R         GE0/0/17     105
GE1/4      00:0b:86:6a:25:40  B:R         GE0/0/18     105

System name
-----------
DellW-3600

DellW-3600


Number of neighbors: 2
```

Display LLDP Neighbor Interface Detail

```
(host) (gigabitethernet "1/3") #show lldp neighbor interface gigabitethernet 1/3 detail
Interface: gigabitethernet1/3, Number of neighbors: 1
------------------------------------------------------------
Chassis id: 24.1.1.253, Management address: 24.1.1.253
Interface description: SW PORT, ID: 04C5A44C3485:P1
Device MAC: 04:c5:a4:4c:34:85
Last Update: Thu Oct  3 17:01:41 2013
Time to live: 180, Expires in: 179 Secs
System capabilities : Bridge,Phone
Enabled capabilities: Bridge,Phone
System name: SEP04C5A44C3485
System description:
Cisco IP Phone 7962G,V10, SCCP42.9-2-1S
Auto negotiation: Supported, Enabled
Autoneg capability:
100Base-X, HD: no, FD: yes
1000Base-T, HD: yes, FD: yes
Media attached unit type: 100BaseTXFD - 2 pair category 5 UTP, full duplex mode (16)
802.3 Power:
PortID:       local 04C5A44C3485:P1
PortDescr:    SW PORT
LLDP-MED:
Device Type:  Communication Device Endpoint (Class III)
Capability:   LLDP-MED capabilities, Network policy, Extended power via MDI/PD, Inventory
LLDP-MED Network Policy for: AppType: 1, Defined: yes
Descr:        Voice
VLAN:         204
Layer 2 Priority: 5
DSCP Value:   46
LLDP-MED Network Policy for: AppType: 2, Defined: yes
Descr:        Voice Signaling
VLAN:         204
Layer 2 Priority: 4
DSCP Value:   32
Extended Power-over-Ethernet:
```

```
Power Type & Source: PD Device
Power Source: unknown
Power Priority: unknown
Power Value:  6300
Inventory:
Hardware Revision: 10
Software Revision: SCCP42.9-2-1S
Firmware Revision: tnp62.8-3-1-21a.bin
Serial Number: FCH1529F57D
Manufacturer: Cisco Systems, Inc.
Model:        CP-7962G
```

Display LLDP Statistics

```
(host)# show lldp statistics
LLDP Statistics
---------------
Interface  Received  Unknow TLVs  Malformed  Transmitted
---------  --------  -----------  ---------  -----------
GE1/3    0         0            0          0
GE1/4    0         0            0          0
```

Display LLDP Statistics Interface

```
(host)# show lldp statistics interface gigabitethernet 1/3

LLDP Statistics
---------------
Interface             Received  Unknow TLVs  Malformed  Transmitted
---------             --------  -----------  ---------  -----------
gigabitethernet1/3    0         0            0          0
```

A *mobility domain* is a group of Dell controllers among which wireless users can roam without losing their IP address. Mobility domains are not tied with the master controller; thus, it is possible for a user to roam between controllers managed by different master controllers, as long as all of the controllers belong to the same mobility domain.

You enable and configure mobility domains only on Dell controllers. No additional software or configuration is required on wireless clients to allow roaming within the domain.

Topics in this chapter include:

- Understanding Dell Mobility Architecture on page 599
- Configuring Mobility Domains on page 600
- Tracking Mobile Users on page 603
- Configuring Advanced Mobility Functions on page 606
- Understanding Bridge Mode Mobility Deployments on page 615
- Enabling Mobility Multicast on page 616

Understanding Dell Mobility Architecture

Dell's layer-3 mobility solution is based on the Mobile IP protocol standard, as described in RFC 3344, IP Mobility Support for IPv4. This standard addresses users who need both network connectivity and mobility within the work environment.

Unlike other layer-3 mobility solutions, a Dell mobility solution does not require that you install mobility software or perform additional configuration on wireless clients. The Dell controllers perform all functions that enable clients to roam within the mobility domain.

In a mobility domain, a *mobile client* is a wireless client that can change its point of attachment from one network to another within the domain. A mobile client receives an IP address (*a home address*) on a *home network*.

A mobile client can detach at any time from its home network and reconnect to a *foreign network* (any network other than the mobile client's home network) within the mobility domain. When a mobile client is connected to a foreign network, it is bound to a *care-of address* that reflects its current point of attachment. A care-of address is the IP address of the Dell controller in the foreign network with which the mobile client is associated.

The *home agent* for the client is the controller at which the client appears for the first time upon joining the mobility domain. The home agent is the single point of contact for the client when the client roams. The *foreign agent* for the client is the controller which handles all Mobile IP communication with the home agent on behalf of the client. Traffic sent to a client's home address is intercepted by the home agent and tunneled for delivery to the client on the foreign network. On the foreign network, the foreign agent delivers the tunneled data to the mobile client.

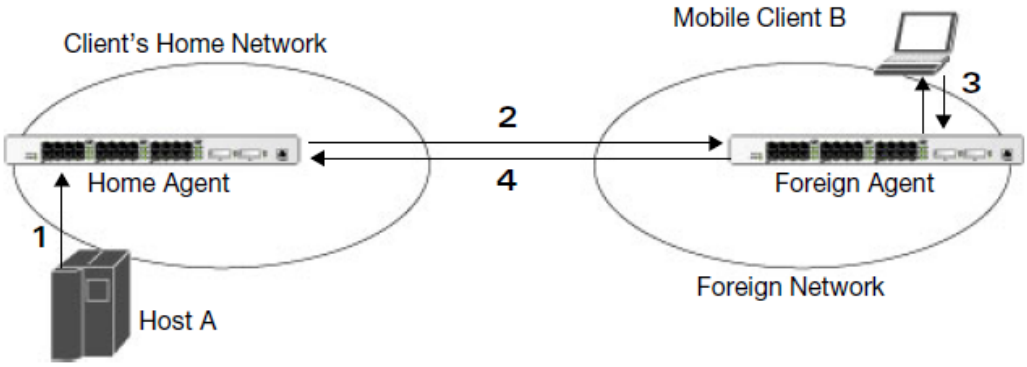
Figure 65 shows the routing of traffic from Host A to Mobile Client B when the client is away from its home network. The client's care-of address is the IP address of the Dell controller in the foreign network.

The numbers in the Figure 65 correspond to the following descriptions:

1. Traffic to Mobile Client B arrives at the client's home network via standard IP routing mechanisms.
2. The traffic is intercepted by the home agent in the client's home network and tunneled to the care-of address in the foreign network.

3. The foreign agent delivers traffic to the mobile client.
4. Traffic sent by Mobile Client B is also tunneled back to the home agent.

Figure 65 *Routing of Traffic to Mobile Client within Mobility Domain*



Configuring Mobility Domains

Before configuring a mobility domain, you should determine the user VLAN(s) for which mobility is required. For example, you may want to allow employees to be able to roam from one subnetwork to another. All controllers that support the VLANs into which employee users can be placed should be part of the same mobility domain.

NOTE

Dell mobility domains are supported only on Dell controllers.

A controller can be part of multiple mobility domains, although it is recommended that a controller belong to only one domain. The controllers in a mobility domain do not need to be managed by the same master controller.

You configure a mobility domain on a master controller; the mobility domain information is pushed to all local controllers that are managed by the same master controller. On each controller, you must specify the *active* domain (the domain to which the controller belongs). If you do not specify the active domain, the controller will be assigned to a predefined "default" domain.

Although you configure a mobility domain on a master controller, the master controller does not need to be a member of the mobility domain. For example, you could set up a mobility domain that contains only local controllers; you still need to configure the mobility domain on the master controller that manages the local controllers. You can also configure a mobility domain that contains multiple master controllers; you need to configure the mobility domain on each master controller.

The basic tasks you need to perform to configure a mobility domain are listed below. The sections following describe each task in further detail. A sample mobility domain configuration is provided in Example Configuration on page 602.

Table 113: *Tasks to Configure a Mobility Domain*

On a master controller:	On all controllers in the mobility domain:
• Configure the mobility domain, including the entries in the home agent table (HAT)	• Enable mobility (disabled by default) • Join a specified mobility domain (not required for "default" mobility domain)

You can enable or disable IP mobility in a virtual AP profile (IP mobility is enabled by default). When you enable IP mobility in a virtual AP profile, the ESSID that is configured for the virtual AP supports layer-3 mobility. If you disable IP mobility for a virtual AP, any clients that associate to the virtual AP will not have mobility service.

Configuring a Mobility Domain

You configure mobility domains on master controllers. All local controllers managed by the master controller share the list of mobility domains configured on the master. Mobility is disabled by default and must be explicitly enabled on all controllers that will support client mobility. Disabling mobility does not delete any mobility related configuration.

In ArubaOS versions before 6.3, the home agent table (HAT) maps a user VLAN IP subnet to potential home agent addresses. Starting from 6.3, when you enable mobility the controller to which the client connects for the first time becomes its home agent. The mobility feature uses the HAT table to locate a potential home agent for each mobile client, and then uses this information to perform home agent discovery. To configure a mobility domain, you must assign a home agent address to at least one controller with direct access to the user VLAN IP subnet. (Some network topologies may require multiple home agents.)

It is recommended that you configure the switch IP address to match the AP's local controller, *or* define the Virtual Router Redundancy Protocol (VRRP) IP address to match the VRRP IP used for controller redundancy. Do not configure both a switch IP address and a VRRP IP address as a home agent address, or multiple home agent discoveries may be sent to the controller.



All user VLANs that are part of a mobility domain must have an IP address that can correctly forward layer-3 broadcast multicast traffic to clients when they are away from the home network.

The mobility domain named "default" is the default active domain for all controllers. If you need only one mobility domain, you can use this default domain. However, you also have the flexibility to create one or more user-defined domains to meet the unique needs of your network topology. Once you assign a controller to a user-defined domain, it automatically leaves the "default" mobility domain. If you want a controller to belong to both the "default" and a user-defined mobility domain at the same time, you must explicitly configure the "default" domain as an active domain for the controller.

Using the WebUI

The following procedure illustrates configuring mobility domain on a master controller:

1. Navigate to the **Configuration** > **Advanced Services** > **IP Mobility** page. Select the **Enable IP Mobility** checkbox.
2. To configure the default mobility domain, select the default domain in the **Mobility Domain** list.

   To create a new mobility domain, enter the name of the domain in **Mobility Domain Name** and click **Add**; the new domain name appears in the **Mobility Domain** list.
3. Select the newly-created domain name. Click **Add** under the Subnet column. Enter the subnetwork, mask, VLAN ID, VRIP, and home agent IP address, and click **Add**. Repeat this step for each HAT entry.
4. Click **Apply**.

Using the CLI

The following command configures mobility domain on a master controller:

```
router mobile
ip mobile domain <name>
   hat <home-agent> description <dscr>
```

To view currently-configured mobility domains in the CLI, use the `show ip mobile domain` command.

Make sure that the ESSID to which the mobile client will connect supports IP mobility. You can disable IP mobility for an ESSID in the virtual AP profile (IP mobility is enabled by default). If you disable IP mobility for a virtual AP, any client that associates to the virtual AP will not have mobility service.

Joining a Mobility Domain

Assigning a controller to a specific mobility domain is the key to defining the roaming area for mobile clients. You should take extra care in planning your mobility domains and survey the user VLANs and controllers to which clients can roam, to ensure that there are no roaming holes.

All controllers are initially part of the "default" mobility domain. If you use the default mobility domain, you do not need to specify this domain as the active domain on a controller. However, once you assign a controller to a user-defined domain, the default mobility domain is no longer an active domain on the controller.

In the WebUI

1. Navigate to the **Configuration** > **Advanced Services** > **IP Mobility** page.
2. In the **Mobility Domain** list, select the mobility domain.
3. Select the **Active** checkbox for the domain.
4. Click **Apply**.

In the CLI

Use the following command to activate a mobility domain:

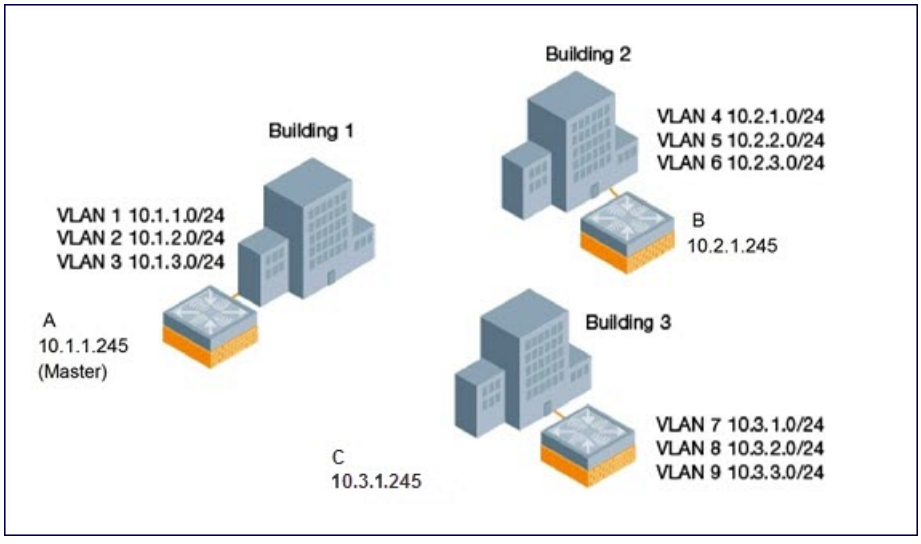
```
ip mobile active-domain <name>
```

To view the active domains in the CLI, use the `show ip mobile active-domains` command on the controller.

Example Configuration

The following example (Figure 66) configures a network in a campus with three buildings. A Dell controller in each building provides network connections for wireless users on several different user VLANs. To allow wireless users to roam from building to building without interrupting ongoing sessions, you configure a mobility domain that includes all user VLANs on the three controllers. You configure the HAT on the master controller only (A in this example). On the local controllers (B and C), you only need to enable mobility and activate the respective domain.

Figure 66 *Example Configuration: Campus-Wide*



This example uses the “default” mobility domain for the campus-wide roaming area. Since all controllers are initially included in the default mobility domain, you do not need to explicitly configure “default” as the active domain on each controller.

Configuring Mobility using the WebUI

On controller A (the master controller):

1. Navigate to the **Configuration** > **Advanced Services** > **IP Mobility** page.
2. Select the **Enable IP Mobility** checkbox.
3. Select the default domain in the Mobility Domain list.
4. Click **Add**.
5. Enter the home agent IP address, and a description for the first entry and click **Add**. Repeat this step for each HAT entry.
6. Click **Apply**.

Table 114: *Example entries*

Home Agent Address or VRIP
10.1.1.245
10.2.1.245
10.3.1.245
10.4.1.245

On controllers B and C:

1. Navigate to the **Configuration** > **Advanced Services** > **IP Mobility** page.
2. Select the **Enable IP Mobility** checkbox.
3. Click **Apply**.

Configuring Mobility using the CLI

On controller A (the master controller):

```
(host)(config) #router mobile
(host)(config) #ip mobile domain default
(host)(mobility-domain) #hat 10.1.1.245 description "Corporate mobile entry"
(host)(mobility-domain) #hat 10.2.1.245 description "Local entry"
(host)(mobility-domain) #hat 10.3.1.245 description "Reserved reentry"
(host)(mobility-domain) #hat 10.4.1.245 description "Sales team"
(host)(mobility-domain) #!
(host)(config)# ip mobile active-domain default
```

On controllers B and C:

```
(host)(config) #router mobile
(host)(config)# ip mobile active-domain default
```

Tracking Mobile Users

This section describes how you can view information about the status of mobile clients in the mobility domain.

Location-related information for users, such as roaming status, AP name, ESSID, BSSID, and physical type are consistent in both the home agent and foreign agent. The username, role, and authentication can be different on the home agent and foreign agent, as explained by the following:

Starting from ArubaOS 6.3, L2 GRE tunnels are automatically established between controllers in mobility domain at the time of boot up. Before ArubaOS 6.3, the tunnels were created only when a client was associated to a controller. Whenever a client connects to a controller in a mobility domain, layer-2 authentication is performed and the station obtains the layer-2 (logon) role. When the client roams to other networks, layer-2 authentication is performed and the client obtains the layer-2 role. If layer-3 authentication is required, this authentication is performed on the client's home agent only. The home agent obtains a new role for the client after layer-3 authentication; this new role appears in the user status on the home agent only. Even if reauthentication occurs after the station moves to a foreign agent, the display on the foreign agent still shows the layer-2 role for the user.

Mobile Client Roaming Status

You can view the list of mobile clients and their roaming status on any controller in the mobility domain:

Viewing Mobile Client Status using the WebUI

Navigate to the **Monitoring** > **controller** > **Clients** page.

Viewing Mobile Client Status using the CLI

Execute the following command to view the status of the mobile client:

```
show ip mobile host
```

Roaming status can be one of the following:

Table 115: *Client Roaming Status*

Roaming Status Type	Description
Home Switch/Home VLAN	This controller is the home agent for a station, and the client is on the VLAN on which it has an IP address.
Mobile IP Visitor	This controller is not the home agent for a client.
Mobile IP Binding (away)	This controller is the home agent for a client that is currently away.
Home Switch/Foreign VLAN	This controller is the home agent for a client, but the client is currently on a different VLAN than its home VLAN (the VLAN from which it acquired its IP address).
Stale	The client does not have connectivity in the mobility domain. Either the controller has received a disassociation message for a client, but has not received an association or registration request for the client from another controller, or a home agent binding for the station has expired before being refreshed by a foreign agent.
No Mobility Service	The controller cannot provide mobility service to this client. The mobile client may lose its IP address if it obtains the address via DHCP and has limited connectivity. The mobile client may be using an IP address that cannot be served, or there may be a roaming hole due to improper configuration.

Viewing User Roaming Status using the CLI

You can view the roaming status of users on any controller in the mobility domain:

```
show user
```

Roaming status can be one of the following:

Table 116: *User Roaming status*

Status Type	Description
Wireless	This client is on its home agent controller and the client is on the VLAN on which it has an IP address.
Visitor	This client is visiting this controller and the controller is not its home agent.
Away	This client is currently away from its home agent controller.
Foreign VLAN	This client is on its home agent controller but the client is currently on a different VLAN than the one on which it has an IP address.
Stale	This should be a temporary state as the client will either recover connectivity or the client's entry is deleted when the stale timer expires.

Viewing specific client information using the CLI

You can use the following CLI command to view the home agent, foreign agent, and roaming status for a specific mobile client:

```
show ip mobile trace <ip-address>|<mac-address>
```

Mobile Client Roaming Locations

You can view information about where a mobile user has been in the mobility domain. This information can only be viewed on the client's home agent.

In the WebUI

1. Navigate to the **Monitoring** > **controller** > **Clients** page.
2. Click **Status**. The mobility state section contains information about the user locations.

In the CLI

Use the following CLI command to view the user locations of a roaming client:

```
show ip mobile trail <ip-address>|<mac-address>
```

HA Discovery on Association

In normal circumstances, a controller performs an HA discovery only when it is aware of the client's IP address which it learns through the ARP or any L3 packet from the client. This limitation of learning the client's IP and then performing the HA discovery is not effective when the client performs an inter switch move silently (does not send any data packet when in power save mode). This behavior is commonly seen with various handheld devices, Wi-Fi phones and so on. This delays HA discovery and eventually results in any loss of downstream traffic that is meant for the mobile client.

When HA discovery on association is triggered, the foreign agent controller to which the client is associated, sends a unicast request to all controllers within the mobility domain to find if any one of the controllers has the IP mobility state information of the client.

With HA discovery on association, a controller can perform a HA discovery as soon as the client is associated. This feature is enabled by default. This option will also poll for all potential HAs.

Setting up Mobility Association using the CLI

Use the following command to configure the mobility association:

```
wlan virtual-ap default ha-disc-onassoc
```

Configuring Advanced Mobility Functions

You can configure various parameters that pertain to mobility functions on a controller in a mobility domain using either the WebUI or the CLI.

In the WebUI

1. Navigate to the **Configuration** > **Advanced Services** > **IP Mobility** page.
2. Select the **Global Parameters** tab.
3. Configure your desired IP mobility settings. Table 117 describes the parameters you can configure on the **Global Parameters** tab.

Table 117: *IP Mobility Configuration Parameters*

Parameter	Description
General	
Encapsulation Supported	This parameter shows the type of encapsulation currently supported on the controller.
Clear Trail Entries	Clear the station location trail table. You can view entries in this table using the `show ip mobile trail` command.
Clear Mobility Counters	Clear counters for IP mobility statistics.
Foreign Agent	
lifetime	Requested lifetime, in seconds, as per RFC 3344, IP Mobility Support for IPv4. Range: 10-65534 seconds Default: 180 seconds
Max. Visitors Allowed	Set a maximum allowed number of active visitors. Range: 0-5000 visitors Default: 5000 visitors
Registration Requests Retransmits	Maximum number of times the foreign agent attempts mobile IP registration message exchanges before giving up. Range: 0-5 attempts Default: 3 attempts
Registration Requests Interval	Retransmission interval, in milliseconds. Range: 100-10000 milliseconds Default: 1000 milliseconds
Home Agent	
Replay	Time difference, in seconds, for timestamp-based replay protection, as described by RFC 3344, IP Mobility Support for IPv4. 0 disables replay. Range: 0-5000 seconds Default: 5000 seconds.

Parameter	Description
Max. Binding Allowed	Maximum number of mobile IP bindings. Note that there is a license-based limit on the number of users and a one user per binding limit in addition to unrelated users. This option is an additional limitation to control the maximum number of roaming users. When the limit is reached, registration requests from the foreign agent fail which causes a mobile client to set a new session on the visited controller, which will become its home controller. Range: 0-300 seconds Default: 7 seconds
Proxy Mobile IP	
Trigger Mobility on Station Association	If enabled, mobility move detection is performed when the client associates with the controller instead of when the client sends packets. This option is enabled by default. Mobility on association can speed up roaming and improve connectivity for devices that do not send many uplink packets out to trigger mobility. The downside to this option is lowered security. An association alone triggers mobility; however, this is irrelevant unless layer-2 security is enforced.
Mobility Trail Logging	Enables logging at the notification level for mobile client moves.
Roaming for Authenticated Stations Only	Allows a client to roam only if has been authenticated. If a client has not been authenticated, no mobility service is offered if it roams to a different VLAN or controller.
Max. Station Mobility Events per Second	Maximum number of mobility events (events that can trigger mobility) handled per second. Mobility events above this threshold are ignored. This helps to control frequent mobility state changes when the client bounces back and forth on APs before settling down. Range: 1-65535 events Default: 25 events
Station Trail Timeout	Specifies the maximum interval, in seconds, an inactive mobility trail is held. Range: 120-86400 seconds Default: 3600 seconds
Station Trail Max. Entries	Specifies the maximum number of entries (client moves) stored in the user mobility trail. Range: 1-100 entries Default: 30 entries.
Mobility Host Entry Hold Time	Number of seconds the mobility state is retained after the loss of connectivity. This allows authentication state and mobility information to be preserved on the home agent controller. The default is 60 seconds but can be safely increased. In many case a station state is deleted without waiting for the stale timeout; user delete from management, foreign agent to foreign agent handoff, and so on. (This is different from the no-service-timeout; no-service-timeout occurs up front, while the stale-timeout begins when mobility service is provided but the connection is disrupted for some reason.)
Mobility Host Entry Lifetime	Time, in seconds, after which mobility service expires. If nothing has changed from the previous state, the client is given another bridge entry but it will have limited connectivity.

Parameter	Description
Revocation	
Retransmits	Maximum number of times the home agent or foreign agent attempts mobile IP registration/revocation message exchanges before giving up. Range: 0-5 retransmissions Default: 3 retransmissions.
Interval	Retransmission interval, in milliseconds. Range: 100-10000 milliseconds Default: 1000 milliseconds

4. Click **Apply**.

In the CLI

To configure foreign agent functionality, use the following command:

```
ip mobile foreign-agent {lifetime <seconds> | max-visitors <number> |
   registrations {interval <msecs> | retransmits <number>}}
```

To configure home agent functionality, use the following command:

```
ip mobile home-agent {max-bindings <number>|replay <seconds>}
```

To configure proxy mobile IP and DHCP functionality, use the following command:

```
ip mobile proxy
   auth-sta-roam-only | event-threshold <number> | log-trail | no-service-timeout <seconds> |
   on-association | stale-timeout <seconds> | trail-length <number> |trail-timeout <seconds>
```

To configure revocation functionality, use the following command:

```
ip mobile revocation {interval <msec>|retransmits <number>
```

To enable packet trace for a given MAC address, use the following command:

```
ip mobile packet-trace <host MAC address>
```

Proxy Mobile IP

The *proxy mobile IP module* in a mobility-enabled controller detects when a mobile client has moved to a foreign network and determines the home agent for a roaming client. The proxy mobile IP module performs the following functions:

- Derives the address of the home agent for a mobile client from the HAT using the mobile client's IP address. If there is more than one possible home agent for a mobile client in the HAT, the proxy mobile IP module uses a discovery mechanism to find the current home agent for the client.
- Detects when a mobile client has moved. Client moves are detected based on ingress port and VLAN changes, and mobility is triggered accordingly. For faster roaming convergence between AP(s) on the same controller, it is recommended that you keep the **on-association** option enabled. This helps trigger mobility as soon as 802.11 association packets are received from the mobile client.

Revocations

A home agent or foreign agent can send a registration revocation message, which revokes registration service for the mobile client. For example, when a mobile client roams from one foreign agent to another, the home agent can send a registration revocation message to the first foreign agent so that the foreign agent can free any resources held for the client.

IPv6 L3 Mobility

ArubaOS supports IPv6 L3 Mobility functionality. The existing L3 mobility solution has been enhanced to support dual stacked (IPv4 and IPv6) and pure IPv6 mobile clients. The IPv6 L3 mobility allows the wireless clients to retain their IPv4 or IPv6 addresses across different VLANs within a controller and between different controllers. In the previous release, the Dell Mobility Controllers supported L3 mobility only for single stacked IPv4 clients.

The following changes in the existing behavior is observed in the Dellcontroller when you enable IPv6 L3 Mobility support :

- The controller throttles and proxies Router Advertisements (RAs) if the router mobile command is enabled.

The following command configures the maximum time allowed between sending unsolicited multicast router advertisements from each interface when RA proxy is enabled:

```
(config)# ipv6 proxy-ra interval <180-1800>
```

The default value for `proxy-ra` interval is 600 seconds. If RA is configured on an external router, but not within the controller, the controller stores the RA in cache and replays the RA from the external server and replays them every proxy-ra interval. If RA is configured in both an external router and in the controller, clients serviced by the controller receive RA only from the controller and not from the external router.

- L3 mobility support for wired and third-party APs are deprecated.
- The HA discovery on association parameter is turned on by default and is not configurable.



By enabling L3 mobility feature, both the solicited RAs and the unsolicited periodic RAs will be converted to L2 unicast and sent to the wireless clients.



It is recommended to reboot the controller when you issue the **no router mobile** command so that mutlicast RAs do not continue to get converted to unicast RAs.

Multicast Mobility

Multicast mobility ensures a client gets an uninterrupted multicast stream while roaming. ArubaOS provides support for a MLD proxy to enable IPv6 multicast mobility. To achieve multicast mobility, the Home Agent (HA) and the Foreign Agent (FA) must be capable of MLD proxying by exchanging the MLD membership information and process MLD messages. ArubaOScontroller supports MLD versions v1 and v2.

Important Points to Remember

- ArubaOS does not support the source-based forwarding functionality of MLDv2.
- The multicast traffic flow stops for few seconds for roaming clients after enabling or disabling the Dynamic Multicast Optimization (DMO) option.

Use the following command to enable MLD proxy in the VLAN:

```
(host)(config)# interface vlan <vlan-id>
(host)(config-subif)# ipv6 mld proxy <gigabitEtherner/fastEhernet> <slot/port>
```

Use the following command to display the interface-specific MLD proxy group information:

```
(host) #show ipv6 mld proxy-group
MLD Proxy Group Table
---------------------
VLAN  Addr                    Group              Num Members
----  ----                    -----              -----------
10    fe80::b:8600:a61:cc5c   ff1e::5            2
10    fe80::b:8600:a61:cc5c   ff02::1:ff9e:dc4c  1
10    fe80::b:8600:a61:cc5c   ff02::1:3          2
```

```
10    fe80::b:8600:a61:cc5c  ff02::1:ff83:d718  1
10    fe80::b:8600:a61:cc5c  ff02::1:ff13:356b  1
10    fe80::b:8600:a61:cc5c  ff02::c            2
Total displayed proxy groups: 6
```

Use the following command to display the MLD proxy mobility database group information for tracking:

```
(host) #show ipv6 mld proxy-mobility-group
MLD MIP Group Table
-------------------
Group Members
----- -------
ff1e::2 1
ff02::1:3 2
ff02::c 1
```

Use the following command to display the statistics of the MLD proxy:

```
(host) #show ipv6 mld proxy-stats
MLD Proxy Statistics(Upstream)
------------------------------
Name     Sent  Received
----     ----  --------
Queries  -     39
Joins    51    112
Leaves   9     0
```

Use the following command to display the MLD proxy mobility multicast statistics:

```
(host)# show ipv6 mld proxy-mobility-stats
MLD Mobility Multicast Statistics
---------------------------------
Name              Sent  Received
----              ----  --------
Joins             -     2
Leaves            -     0
Intra-move        -     1
Inter-move        -     0
Client-away       -     0
Back-home         -     0
Query-db          -     0
Query-foreign-db  -     0
Query-home-db     -     0
Add-visitor       -     0
Replies           0     -
```

The following command displays the discovery count table that keeps track of per client home agent discovery:

```
(host) #show datapath mobility discovery-table
Datapath Mobility Discovery Count Table
-----------------------------------------------
Index     Valid    Version   Retry#    No-Response    Ack      Mac                Vlan
-------   ------   -------   ------    -----------    ------   --------------     -----
1         1         2        1         a              0        10:78:D2:FA:7D:38  74
```

The following command displays the datapath HA table information:

```
(host) #show datapath mobility home-agent-table
Datapath Mobility Home Agent Table
----------------------------------
Switch IP
---------------
10.16.19.14
10.16.19.140
```

The following command displays the mobility multicast-group table that floods the multicast RA traffic to the roaming clients:

```
(host) # show datapath mobility mcast-table
Datapath Mobility Multicast Table
---------------------------------
GRE Tunnel  HomeVlan  McastGroup  Members
----------  --------  ----------  -------
0x10009       501           0  1
```

The following commands displays the statistics of the datapath mobility:

```
(host) #show datapath mobility stats
Datapath Mobility Stats
Mcast group entry alloc errors          : 0
Frames flooded over MMG (@HA)           : 0
Frames subjected to MMG (@FA)           : 0
Frames sent to roamed clients           : 0
HA Discovery failure to notify NACK     : 0
HA Discovery invalid DCT                : 0
HA Discovery DCT allocation failed      : 0
HA Discovery Probes sent                : 0
HA Discovery NULL bridge entry in DCT   : 0
HA Discovery failed to start            : 0
HA Discovery successfully started       : 0
HAT insert failure                      : 0
HAT insert success                      : 0
HAT delete failure                      : 0
HAT delete success                      : 0
```

The following command displays the mobility multicast VLAN table information:

```
(host) #show ip mobile multicast-vlan-table
Mobility Multicast Vlan Table
-----------------------------
Client MAC           Home vlan  Current vlan
----------           ---------  ------------
40:2C:F4:36:16:07   501        501
```

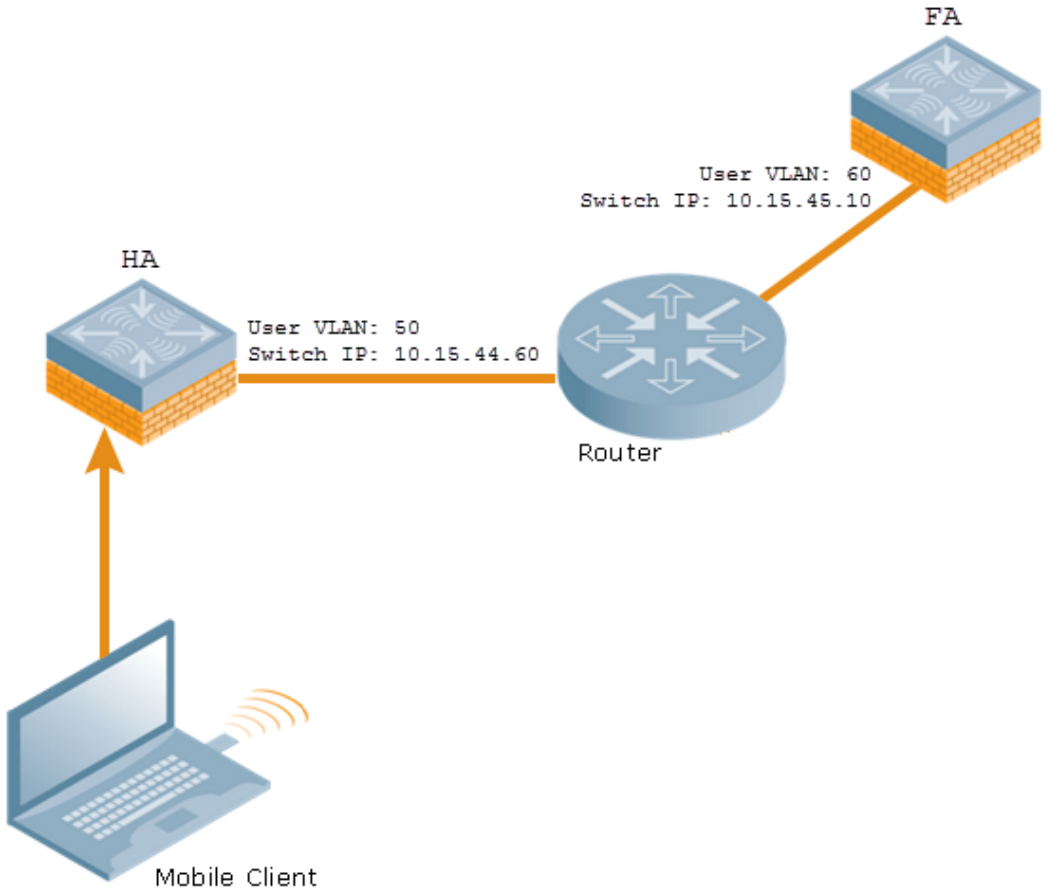
The outputs of the following commands are enhanced to support IPv6 L3 mobility:

- `show ip mobile host`
- `show ip mobile trace`
- `show ip mobile remote`
- `show ip mobile binding`
- `show ip mobile visitor`
- `show ip mobile trail`
- `show ip mobile packet-trace`
- `clear ip mobile trail <IPv6_addr>`
- `show ip mobile traffic`
- `show ip mobile global`
- `show ip mobile hat`
- `show ip mobile domain`
- `ip mobile domain <name> hat <home-agent> description <dscr>`

Example Configuration

The following figure provides information on how a client moves from one controller to another, when you enable IPv6 L3 mobility feature:

Figure 67 *Sample IPv6 L3 Mobility Configuration*



The following commands displays the initial configuration on HA and FA:

```
(host-HA) #show ip mobile domain
Mobility Domains:, 2 domain(s)
------------------------------
Domain name default
Home Agent Table
Domain name 6.3mobility
Home Agent Table
Home Agent      Description
--------------- ----------------
10.15.45.10
10.15.44.60

(host-FA) #show ip mobile domain
Mobility Domains:, 2 domain(s)
------------------------------
Domain name default
Home Agent Table
Domain name 6.3mobility
Home Agent Table
Home Agent      Description
--------------- ----------------
10.15.45.10
10.15.44.60
```

The following commands displays information on the client association to HA:

```
(host-HA) #show user
Users
-----
IP                              MAC             Name     Role           Age(d:h:m)  Auth  VPN
link  AP name             Roaming   Essid/Bssid/Phy                       Profile  Forward
mode  Type  Host Name
----------                      ------------       ------    ----           ----------  ----  -
-------  -------             -------   ---------------                       -------  --------
----  ----  -

50.50.50.11                     24:77:03:9e:dc:4c            authenticated  00:00:00
        AP124-B11550-Jibin  Wireless  mobility-test/00:1a:1e:82:b3:10/a-HT  default  tunnel

2001:5000::2677:3ff:fe9e:dc4c  24:77:03:9e:dc:4c            authenticated  00:00:00
        AP124-B11550-Jibin  Wireless  mobility-test/00:1a:1e:82:b3:10/a-HT  default  tunnel
fe80::2677:3ff:fe9e:dc4c       24:77:03:9e:dc:4c            authenticated  00:00:00
        AP124-B11550-Jibin  Wireless  mobility-test/00:1a:1e:82:b3:10/a-HT  default  tunnel

(host-HA) #show ip mobile host
Mobile Host List, 1 host(s)
---------------------------
24:77:03:9e:dc:4c
IPv4: 50.50.50.11
IPv6: fe80::2677:3ff:fe9e:dc4c, 2001:5000::2677:3ff:fe9e:dc4c
Roaming Status: Home Switch/Home VLAN, Service time 0 days 00:00:57
Home VLAN 50

(host-HA) #show datapath bridge table 24:77:03:9e:dc:4c
Datapath Bridge Table Entries
-----------------------------
Flags: P - Permanent, D - Deny, R - Roamed Client, M - Mobile, X - Xsec, A - Auth, O - Outer
VLAN, T - Trusted
MAC          VLAN  Assigned VLAN  QinQ VLAN  Destination  Flags
-----------------  ----  -------------  ---------  -----------  -----
24:77:03:9E:DC:4C  50    50             0          tunnel 17    PM

(host-HA) #show datapath station
Datapath Station Table Entries
------------------------------
Flags: W - WEP, T - TKIP, A - AESCCM, M - WMM N - .11n client
S - AMSDU, G - AESGCM, R - DATA READY, I - INACTIVE, r - ROAMED
MAC               BSSID       VLAN Bad Decrypts Bad Encrypts Cpu Qsz      RSN cap Aid  HomeVlan
Flags
----------------- ----------------- ---- ------------ ------------ --- -------- ------- ---- -
------- -----
24:77:03:9E:DC:4C 00:1A:1E:82:B3:10   50            0            0 8    0 0 0 0    0000 0001
    50 MN
```

The following commands displays status of the client roaming to FA:

```
(host-FA) #show ap association
Association Table
-----------------
Name      bssid             mac               auth  assoc  aid  l-int  essid          vlan-
id  tunnel-id  phy            assoc. time  num assoc  Flags  Band steer moves (T/S)
----      -----             ---               ----  -----  ---  -----  -----          ------
-  ---------  ---            -----------  ---------  -----  ----------------------
Ap_local  6c:f3:7f:3a:ba:d8  24:77:03:9e:dc:4c  y     y      1    100    mobility-test  60
  0x1000f    a-HT-40sgi-2ss  3m:20s       1          WA     0/0
```

```
Num Clients:1
(host-FA) #show us
Users
-----
IP                                MAC            Name     Role                    Age(d:h:m)
Auth  VPN link  AP name   Roaming  Essid/Bssid/Phy                        Profile  Forward mode
 Type   Host Name
----------                        ------------       ------    ----                    ----------
----  --------  -------   -------  ---------------                        -------  ------------
 ----   --
50.50.50.11                     24:77:03:9e:dc:4c            sys_mip_role_649130_9  00:00:03
               Ap_local  Visitor  mobility-test/6c:f3:7f:3a:ba:d8/a-HT  default  tunnel
 Win 7
2001:5000::2677:3ff:fe9e:dc4c  24:77:03:9e:dc:4c            sys_mip_role_649130_9  00:00:03
               Ap_local  Visitor  mobility-test/6c:f3:7f:3a:ba:d8/a-HT  default  tunnel
 Win 7
User Entries: 2/2
Curr/Cum Alloc:1/7 Free:1/6 Dyn:2 AllocErr:0 FreeErr:0
(host-FA) #show ip mobile host
Mobile Host List, 1 host(s)
---------------------------
24:77:03:9e:dc:4c
IPv4: 50.50.50.11
IPv6: 2001:5000::2677:3ff:fe9e:dc4c
Roaming Status: Mobile IP Visitor, Service time 0 days 00:03:33
Home VLAN 50, visiting local VLAN 60
(host-FA) #show datapath bridge table 24:77:03:9e:dc:4c
Datapath Bridge Table Entries
-----------------------------
Flags: P - Permanent, D - Deny, R - Roamed Client, M - Mobile, X - Xsec, A - Auth, O - Outer
VLAN, T - Trusted
MAC          VLAN  Assigned VLAN  QinQ VLAN  Destination  Flags
-----------------  ----  -------------  ---------  -----------  -----
24:77:03:9E:DC:4C  4095  60             0          tunnel 15    PMR
24:77:03:9E:DC:4C  60    60             0          tunnel 15    PM

(host-FA) #show datapath station
Datapath Station Table Entries
------------------------------
Flags: W - WEP, T - TKIP, A - AESCCM, M - WMM N - .11n client
S - AMSDU, G - AESGCM, R - DATA READY, I - INACTIVE, r - ROAMED
MAC               BSSID       VLAN Bad Decrypts Bad Encrypts Cpu Qsz      RSN cap Aid  HomeVlan
Flags
----------------- ----------------- ---- ------------ ------------ --- -------- ------- ---- -
------- -----
24:77:03:9E:DC:4C 6C:F3:7F:3A:BA:D8   60            0            0 7    0 0 0 0    0000 0001
    50 MNr

(host-FA) #show ip mobile visitor
Foreign Agent Visitor list, 1 host(s)
-------------------------------------
24:77:03:9e:dc:4c
IPv4: 50.50.50.11
IPv6: 2001:5000::2677:3ff:fe9e:dc4c
HA Addr 10.15.44.60, Registration id D51BA8BC:856865FC
Lifetime granted 00:00:40 (40), remaining 00:00:36
Tunnel id  9, src 10.15.44.10 dest 10.15.44.60, reverse-allowed
```

The following command displays the status of the client on HA after roaming:

```
(host-HA) #show user
Users
```

```
-----
IP                                MAC             Name      Role            Age(d:h:m)   Auth   VPN
link   AP name    Roaming   Essid/Bssid/Phy                        Profile   Forward mode   Type
Host Name
----------                      ------------       ------    ----           ----------   ----   -
-------   -------    -------   ---------------                     -------   ------------   ----
---------
50.50.50.11                     24:77:03:9e:dc:4c            authenticated   00:00:08
        Ap_local   Away     mobility-test/6c:f3:7f:3a:ba:d8/a-HT   default   tunnel
2001:5000::2677:3ff:fe9e:dc4c   24:77:03:9e:dc:4c            authenticated   00:00:08
        Ap_local   Away     mobility-test/6c:f3:7f:3a:ba:d8/a-HT   default   tunnel
User Entries: 2/2
Curr/Cum Alloc:1/16 Free:1/15 Dyn:2 AllocErr:0 FreeErr:0

(host-HA) #show ip mobile host
Mobile Host List, 1 host(s)
---------------------------
24:77:03:9e:dc:4c
IPv4: 50.50.50.11
IPv6: 2001:5000::2677:3ff:fe9e:dc4c
Roaming Status: Mobile IP Binding (Away), Service time 0 days 00:08:20
Home VLAN 50

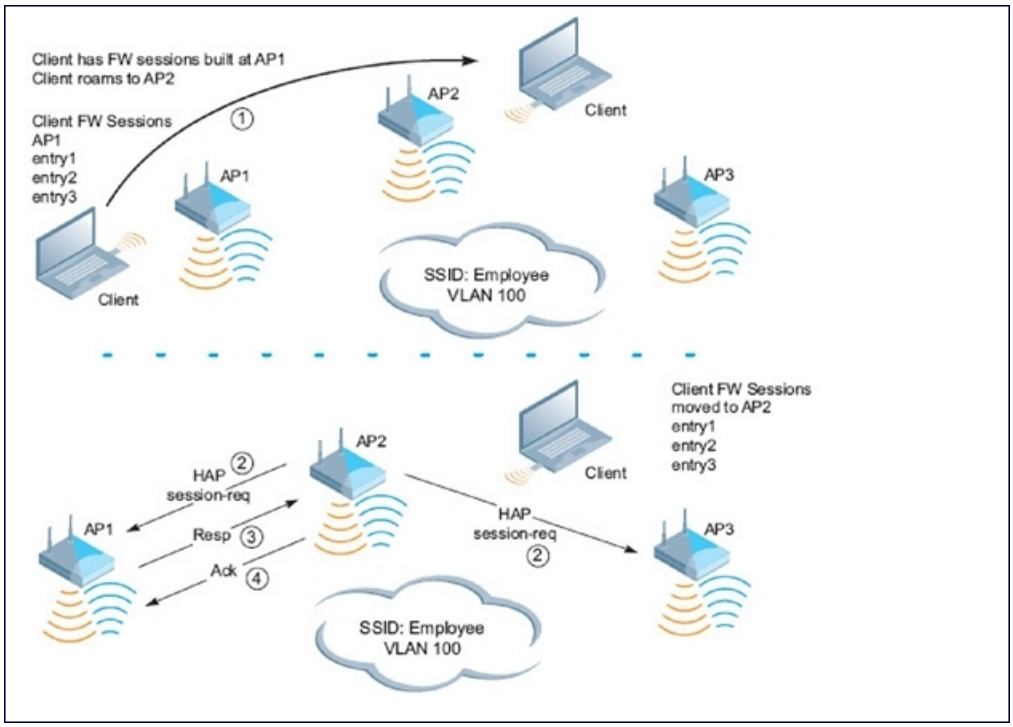
(host-HA) #show datapath bridge table 24:77:03:9e:dc:4c
Datapath Bridge Table Entries
-----------------------------
Flags: P - Permanent, D - Deny, R - Roamed Client, M - Mobile, X - Xsec, A - Auth, O - Outer
VLAN, T - Trusted
MAC          VLAN  Assigned VLAN  QinQ VLAN  Destination  Flags
-----------------  ----  -------------  ---------  -----------  -----
24:77:03:9E:DC:4C  4095  50             0          tunnel 9     PMT
24:77:03:9E:DC:4C  50    50             0          tunnel 9     PMTR
(host-HA) #show ip mobile binding
Home Agent Binding list, 1 host(s)
----------------------------------
24:77:03:9e:dc:4c
IPv6: 2001:5000::2677:3ff:fe9e:dc4c
FA Care-of Addr 10.15.44.10, Src Addr 10.15.44.10, HAT HA Addr 10.15.44.60
FA Visiting VLAN 60
Lifetime granted 00:00:40 (40), remaining 00:00:23
Flags T, Registration id D51BA8BC:856865FC
Tunnel id  9, src 10.15.44.60 dest 10.15.44.10, reverse-allowed
```

Understanding Bridge Mode Mobility Deployments

In bridge mode deployments, it is possible to deploy more than one AP in a single location. Therefore, APs in bridge forwarding mode support firewall session synchronization, which allows clients to retain their current session and IP address as they roam between different bridge mode APs on the same layer-2 network.

The bridge mode mobility feature facilitates client mobility on up to 32 layer-2 connected APs by allowing the APs to communicate and share the user state as the user roams from AP to AP. This mechanism is always enabled when an AP is set to bridge mode, and it requires that all APs be on the same Layer 2 segment where roaming will occur.

Figure 68 *Bridge Mode Mobility*



The roaming process occurs as follows:

1. A client begins to roam from AP1 and starts an association with AP2.
2. AP2 sends a broadcast message to all APs on the local layer-2 network, asking if any other AP has a current session state for the roaming client.
3. Only AP1 responds to the broadcast, and sends the current session table of the client.
4. AP2 acknowledges receipt of the session table.
5. AP1 deletes the session state of the client.
6. Roaming is complete.

Enabling Mobility Multicast

Internet Protocol (IP) multicast is a network addressing method used to simultaneously deliver a single stream of information from one sender to multiple clients on a network. Unlike broadcast traffic, which is meant for all hosts in a single domain, multicast traffic is sent only to those specific hosts who are configured to receive such traffic. Clients who want to receive multicast traffic can join a multicast group via IGMP messages. Upstream routers use IGMP message information to compute multicast routing tables and determine the outgoing interfaces for each multicast group stream.

In ArubaOS 3.3.x and earlier, when a mobile client moved away from its local network and associated with a VLAN on a foreign controller (or a foreign VLAN on its own controller), the client's multicast membership information would not be available at its new destination, and multicast traffic from the client could be interrupted. ArubaOS 3.4 and later supports mobility multicast enhancements that provide uninterrupted streaming of multicast traffic, regardless of a client's location.

Working with Proxy IGMP and Proxy Remote Subscription

The mobility controller is always aware of the client's location, so the controller can join multicast group(s) on behalf of that mobile client. This feature, called Proxy IGMP, allows the controller to join a multicast group and

suppresses the client's IGMP control messages to the upstream multicast router. (The client's IGMP control messages will, however, still be used by controller to maintain a multicast forwarding table.) The multicast IGMP traffic originating from the client will instead be sent from the controller's incoming VLAN interface IP.

The IGMP proxy feature includes both a host implementation and a router implementation. An upstream router sees a controller running IGMP proxy as a host; a client attached to the controller sees the controller as router. When you enable Proxy IGMP, all multicast clients associated with the controller are hidden from the upstream multicast device or router.



The newer IGMP proxy feature and the older IGMP snooping feature cannot be enabled at the same time, as both features add membership information to multicast group table. For most multicast deployments, you should enable the IGMP Proxy feature on all VLAN interfaces to manage all the multicast membership requirements on the controller. If IGMP snooping is configured on some interfaces, there is a greater chance that multicast information transfers may be interrupted.

IGMP proxy must be enabled or disabled on each individual interface. To use the IGMP proxy, ensure that the VLANs on the controllers are extended to the upstream router. Enabling IGMP proxy enables IGMP on the interface and sets the querier to the controller itself. You must identify the controller port from which the controller sends proxy join information to the upstream router, and identify the upstream router by upstream port so the controller can dynamically update the upstream multicast router information.

IGMPv3 Support

ArubaOS 6.4 supports IGMPv3 functionality that makes Dell controllers aware of the Source Specific Multicast (SSM) and is used to optimize bandwidth of the network. The SSM functionality is an extension of IP multicast where the datagram traffic is forwarded to receivers from only those multicast sources to which the receivers have explicitly joined. By default, the multicast group range of 232.0.0.0 through 232.255.255.255 (232/8) is reserved for SSM by IANA (Internet Assigned Numbers Authority).

The IGMPv3 snooping functionality is configured at the edge of the network. The devices that support IGMP snooping listen for the IGMP messages that the host sent to join an IP multicast group. These devices record details of all the hosts and also about the IP multicast group in which a particular host has joined. These devices forward IP multicast traffic to the hosts that have joined the specific multicast group.



The IGMP proxy and IGMP snooping functionalities cannot be enabled on the same VLAN simultaneously.

Configuring SSM Range

You can configure the SSM range by using the CLI and WebUI.

Using the CLI

Use the following command:

```
(host)(config) # ip igmp
(host)(config-ip)# ssm-range <startip> <maskip>
```

Using the WebUI

1. Navigate to **Configuration > Network > IP > Multicast Routing** page of the WebUI.
2. In the **IGMP** tab, enter values for SSM Range in the **SSM Range Start-IP** and **SSM Range Mask-IP** text boxes.
3. Click **Apply**.

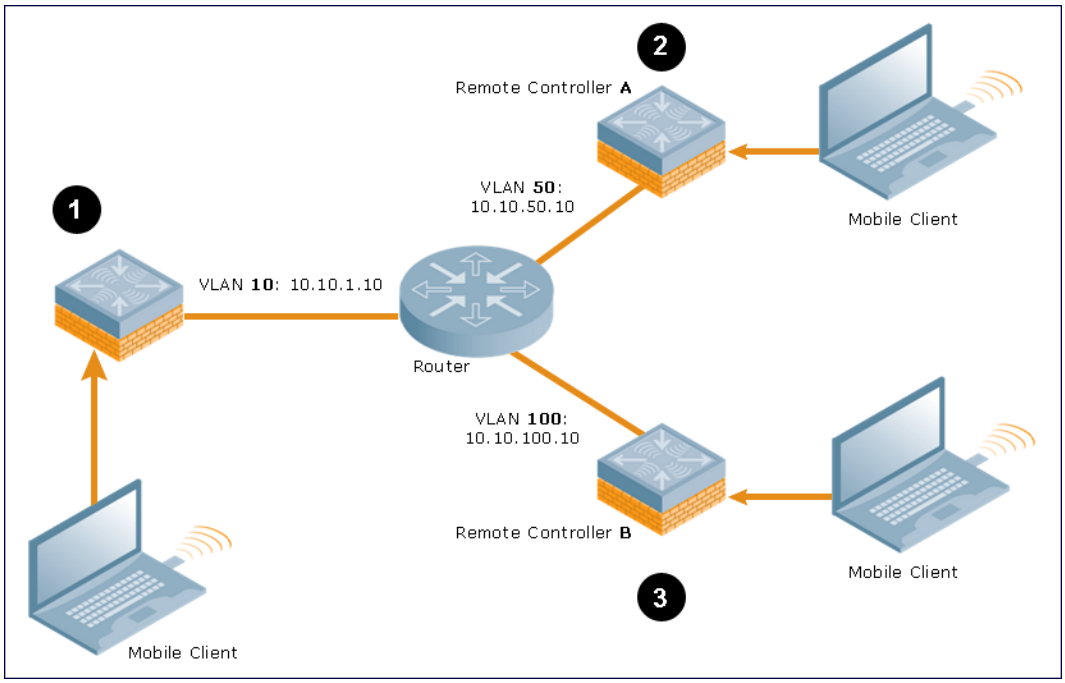
NOTE

The proxy operation will be downgraded to IGMPv2 if any lower version clients are present and reverts back to v3 mode if the controller finds no lower version client joins (reports) for a specified interval of time. In the downgraded proxy operation the SSM semantics is not applicable for the particular VLAN.

Working with Inter Controller Mobility

When a client moves from one controller to another, multicast traffic migrates as follows:

Figure 69 *Inter-Controller Mobility*



1. The local controller uses its VLAN 10 IP address to join multicast group1 on behalf of a mobile client.
2. The mobile client leaves its local controller and roams to VLAN 50 remote controller A.

   Remote controller A locates the mobile client's local controller and learns about the client's multicast groups. Remote controller A then joins group1 on behalf the mobile client, using its VLAN 50 source IP. Upstream multicast traffic from the roaming client is sent to the local controller over an L2 GRE tunnel. The remote controller will receive downstream multicast traffic and send it to the mobile client.

NOTE

The L2-GRE Tunnel implementation of the IP mobility functionality is supported only on ArubaOS versions 6.2 or later, and is not backward compatible with the earlier implementation. ArubaOS supports only v4 mobility and does not support IPv6 L3 mobility.

   Meanwhile, the local controller checks to see if other local clients require group1 traffic. If no other clients are interested in group1, then the local controller will leave that group. If there are other clients using that group, the controller will continue its group1 membership.
3. Now the mobile client leaves remote controller A and roams to VLAN 100 on remote controller B. Remote controller B locates the mobile client's local controller and learns about the client's multicast groups. Remote controller B then joins group1 on behalf the roaming mobile client 1, using its VLAN 100 IP address.

   Both the local controller and remote controller A will verify if any of their other clients require group1 traffic. If none of their other clients require group1, then that controller will leave the group. (If the local controller leaves the group, it will also notify remote controller A.) If either controller has other clients using that group, that controller it will continue its group1 membership.

Configuring Mobility Multicast

In the WebUI

To configure the mobility multicast feature using the controller WebUI:

1. Navigate to the **Configuration** > **Network** > **IP** window.
2. Click **Edit** by the VLAN interface for which you want to configure mobility multicast. The **Edit VLAN** window opens.
3. Select **Enable IGMP** to enable the router to discover the presence of multicast listeners on directly-attached links.
4. Select **Snooping** to save bandwidth and limit the sending of multicast frames to only those nodes that need to receive them.
5. Select the **Interface** checkbox, then click the **Proxy** drop-down list and select the controller interface, port and slot for which you want to enable proxy IGMP.
6. Click **Apply**.
7. (Optional) Repeat steps 1-6 above to configure mobility multicast for another VLAN interface.

In the CLI

The following command enables IGMP and/or IGMP snooping on this interface, or configures a VLAN interface for uninterrupted streaming of multicast traffic:

```
interface vlan <vlan>
   ip igmp proxy [{fastethernet|gigabitethernet} <slot>/<port>]|[snooping]
```

Table 118: *Command Syntax*

Parameter	Description
`fastethernet`	Enable IGMP proxy on the FastEthernet (IEEE 802.3) interface
`gigabitethernet`	Enable IGMP proxy on the GigabitEthernet (IEEE 802.3) interface

`<slot>/<port>`
`<slot>/<module>/<port>`
`(W-7200only)`
	Any command that references a Fast Ethernet or Gigabit Ethernet interface requires that you specify the corresponding port on the controller in the format <slot>/<port>.
	<slot> is always 1, except when referring to interfaces on the W-6000 controller(slots 0-3).
	The **<port>** parameter refers to the network interfaces that are embedded in the front panel of the W-3000 Series controller, or a W-6000M3 controller module. Port numbers start at 0 from the left-most position.
	The W-7200 Series controllers use a <slot>/<module>/<port> port numbering scheme. <slot> and <module> will always be 0 on the W-7200 Series.

`snooping`	Enable IGMP snooping. The IGMP protocol enables an router to discover the presence of multicast listeners on directly-attached links. Enable IGMP snooping to limit the sending of multicast frames to only those nodes that need to receive them.

Example

The following example configures IGMP proxy for VLAN 2. IGMP reports from the controller would be sent to the upstream router on fastethernet port 1/3:

```
conf# interface vlan 2
   conf-subif# ip igmp proxy fastethernet 1/3
```

User-Identification (User-ID) feature of the Palo Alto Networks (PAN) firewall allows network administrators to configure and enforce firewall policies based on user and user groups. User-ID identifies the user on the network based on the IP address of the device which the user is logged into. Additionally, firewall policy can be applied based on the type of device the user is using to connect to the network. Since the Dell controller maintains the network and user information of the clients on the network, it is the best source to provide the information for the User-ID feature on the PAN firewall.



PAN firewall integration with ArubaOS requires PAN-OS 5.0 or later

This feature introduces the following interactions with PAN firewall servers:

- Send logon events to PAN firewall for the client with its IP address user name, device type, when classified.
- Send logout events to PAN firewalls for the client with its IP address.

The following must be configured on the PAN Firewall:

- An Admin account must be created on the PAN firewall to allow the controller to send data to the PAN firewall. This account must match the account added in the PAN profile on the controller. The built-in Admin account can be used for this purpose but that is not recommended. It is better to create a new Admin account used solely for the purpose of communications between the controller and PAN firewall.
- Preconfiguration of PAN Host Information Profile (HIP) objects and HIP-profiles on the PAN Firewall to support a device-type based policy.

To enable these features, the following must be configured on the controller:

- System wide PAN profile must be properly configured and made active on the controller.
- The **pan-integration** knob at the AAA profile which the client is associated with must be enabled.
- For VPN clients, the **pan-integration** knob in the VPN authentication profile which the client is associated with must be enabled.
- For VIA clients, the **pan-integration** knob in the VIA authentication profile which the client is associated with must be enabled.

Limitations

Keep the following limitations when in mind when configuring PAN Firewall Integration:

- PAN Firewall Integration does not support bridge forwarding mode.
- The W-600 Series controller does not support PAN Firewall integration.

Preconfiguration on the PAN Firewall

Before PAN Firewall configuration is completed on the controller, some configuration must be completed on the PAN Firewall.

User-ID Support

The administrator must configure firewall policies based on user-name and/or user-group. Additionally, correct configuration of connection to directory servers is needed for user-group based policies on the PAN firewall.

Device-Type Based Policy Support

The controller supports a limited number of device types. The identified device type associated with each IP user will be sent to the PAN in the **client-version** field with the **host-info** category of the HIP-report. PAN administrators must create these HIP objects, which filter the HIP-reports sent from the controller to support device-type based firewall policies.

Table 119 lists the HIP objects with a specified **Is Value** in the **Client Version** field, which must be preconfigured on the PAN firewall.

Table 119: *HIP Objects*

Client Version Is Value
Android
Apple
AppleTV
BlackBerry
Chrome OS
iPad
iPhone
iPod
Kindle
Linux
Nintendo
Nintendo 3DS
Nintento Wii
Nook
OS X
PlayStation
PS Vita
PS3

Client Version Is Value
PSP
RIM Tablet
Roku
Symbian
webOS
Win 7
Win 8
Win 95
Win 98
Win 2000
Win CE
Win ME
Win NT
Win Server
Win Vista
Win XP
Windows
Windows Mobile
Windows Phone 7

Configuring PAN Firewall Integration

A PAN profile must be created on the controller. Multiple PAN profiles can be configured and saved on the controller but only one profile can be active at a time. These profiles can be configured and applied using the ArubaOS WebUI or CLI.

Creating PAN Profiles

The first step in configuring PAN firewall integration is to create PAN Profiles. This profile provides the controller with the information required for connecting to and interacting with the specified PAN firewall. The PAN profile can be created using the WebUI or CLI.

NOTE

This configuration is done and available on the master controller only. The configuration will be pushed to all connected local controllers.

Using the WebUI

To configure a new PAN profile, complete the following steps:

1. Navigate to **Configuration > Advanced Services > All Profiles > Other Profiles > Palo Alto Networks Servers**.
2. Type the name of the PAN profile and click **Add**.
3. Click on the name of the name PAN profile to open the **Profile Details** window.
4. Enter the **Host (IP address or hostname)** of the PAN firewall
5. Enter the **Port (1 – 65535)** of the PAN Firewall.
6. Enter the **Username** of the PAN firewall. The user name is between 1 and 255 bytes in length. The username must match the Admin account previously created on the PAN firewall.
7. Enter the **Password** of the username in PAN Firewall. The password is between 6 and 100 bytes in length. The password must match the Admin account previously created on the PAN firewall.
8. Re-enter the **Password** entered in the previous step.
9. Click **Add**.
10. Click **Apply**.

NOTE

Up to twenty (20) PAN firewalls are supported.

Table 120: *PAN Profile Parameters*

Parameter	Description
Host (IP or hostname)	The hostname or IP address of the PAN firewall.
Port (1 - 65535)	The port number of the PAN firewall.
Username	The username in the PAN firewall (1 - 255 bytes in length).
Password	Enter the password of the PAN firewall.
Retype Password	Retype the password of the PAN firewall.

Using the CLI

```
(host)(config) #pan profile <profile-name>
   firewall host <host> port <port> username <username> passwd <password>
```

Activating a PAN Profile

Once a PAN profile has been created, the profile must be activated. Select profile you want to activate from the list of configured profiles.

NOTE

This configuration must be completed on each local controller.

Using the WebUI

To apply a PAN profile, complete the following steps:

1. Navigate to **Configuration > Advanced Services > All Profiles > Other Profiles > Palo Alto Networks Active**.
2. Select **Active Palo Alto Networks**. To the right of this link, the name of the active profile is displayed.
3. Another configured profile can be selected from the **Active Palo Alto Networks Profile >** drop-down menu. Additionally, a new profile can be configured by selecting **--NEW--** and completing the configuration details.
4. Once a profile is selected from the drop-down menu or a new profile is created, click **Apply**.

Using the CLI

```
(host)(config) #pan active-profile
  profile <profile- name>
```

Enabling PAN Firewall Integration

PAN firewall integration must be enabled on the AAA profile that the client is associated with.

Using the WebUI

To enable PAN firewall integration in the AAA profile:

1. Navigate to the **Configuration > Security > Authentication > AAA Profiles** page.
2. In the **AAA Profiles Summary**, select the desired profile.
3. Check the **PAN Firewalls Integration** check box.
4. Click **Apply**.

Using the CLI

```
(host)(config) #aaa profile <aaa profile-name>
  pan-integration
```

Enabling PAN Firewall Integration for VIA Clients

For VIA clients, PAN firewall integration must be enabled on the VIA authentication profile that the client is associated with.

Using the WebUI

To enable PAN firewall integration for VIA clients:

1. Navigate to the **Security > Authentication > L3 Authentication** page.
2. In the profiles list on the left, click **VIA Authentication** and select the desired profile.
3. Check the **PAN Firewalls Integration** check box.
4. Click **Apply**.

Using the CLI

```
(host)(config) #aaa authentication via auth-profile <profile-name>
  pan-integration
```

Enabling PAN Firewall Integration for VPN Clients

For VPN clients, PAN firewall integration must be enabled on the VPN authentication profile that the client is associated with.

Using the WebUI

To enable PAN firewall integration for VPN clients:

1. Navigate to the **Security > Authentication > L3 Authentication** page.
2. In the profiles list on the left, click **VPN Authentication** and select the default profile.
3. Check the **PAN Firewalls Integration** check box.
4. Click **Apply**.

Using the CLI

```
(host)(config) #aaa authentication vpn default
  pan-integration
```

In many deployment scenarios, an external firewall is situated between Dell devices. This chapter describes the network ports that need to be configured on the external firewall to allow proper operation of the Dell network. You can also use this information to configure session ACLs to apply to physical ports on the controller for enhanced security. However, that this chapter does not describe requirements for allowing specific types of user traffic on the network.



A controller uses both its loopback address and VLAN addresses for communications with other network elements. If the firewall uses host-specific ACLS, those ACLs must specify all IP addresses used on the controller.

Topics in this chapter include:

- Understanding Firewall Port Configuration Among Dell Devices on page 626
- Enabling Network Access on page 627
- Ports Used for Virtual Internet Access (VIA) on page 627
- Configuring Ports to Allow Other Traffic Types on page 627

Understanding Firewall Port Configuration Among Dell Devices

This section describes the network ports that need to be configured on the firewall to allow proper operation of the network.

Communication Between Controllers

Configure the following ports to enable communication between any two controllers:

- IPSec (UDP ports 500 and 4500) and ESP (protocol 50). PAPI between a master and a local controller is encapsulated in IPSec.
- IP-IP (protocol 94) and UDP port 443 if Layer-3 mobility is enabled.
- GRE (protocol 47) if tunneling guest traffic over GRE to DMZ controller.
- IKE (UDP 500).
- ESP (protocol 50).
- NAT-T (UDP 4500).

Communication Between APs and the Controller

APs use Trivial File Transfer Protocol (TFTP) during their initial boot to grab their software image and configuration from the controller. After the initial boot, the APs use FTP to grab their software images and configurations from the controller. In many deployment scenarios, an external firewall is situated between various Dell devices.

Configure the following ports to enable communication between an AP and the controller:

- PAPI (UDP port 8211). If the AP uses DNS to discover the LMS controller, the AP first attempts to connect to the master controller. (Also allow DNS (UDP port 53) traffic from the AP to the DNS server.)
- PAPI (UDP port 8211). All APs running as Air Monitors (AMs) require a permanent PAPI connection to the master controller.
- FTP (TCP port 21).

- TFTP (UDP port 69) all APs, if there is no local image on the AP (for example, a new AP) the AP will use TFTP to retrieve the initial image.
- SYSLOG (UDP port 514).
- PAPI (UDP port 8211).
- GRE (protocol 47).

Communication Between Remote APs and the Controller

Configure the following ports to enable communication between a Remote AP (IPSec) and a controller:

- NAT-T (UDP port 4500).
- TFTP (UDP port 69) .



TFTP is not needed for normal operation. If the remote AP loses its local image for any reason, it will use TFTP to download the latest image.

Enabling Network Access

This section describes the network ports that need to be configured on the firewall to manage the Dell network.

For WebUI access between the network administrator's computer (running a Web browser) and a controller:

- HTTP (TCP ports 80 and 8888) or HTTPS (TCP ports 443 and 4343).
- SSH (TCP port 22 or TELNET (TCP port 23).

Ports Used for Virtual Internet Access (VIA)

The following ports are used with Dell VIA.

- For the reachability/trusted network check use port 443
- For the IPSec connection use port 4500
- To allow ISAKMP use port 500

Configuring Ports to Allow Other Traffic Types

This section describes the network ports that need to be configured on the firewall to allow other types of traffic in the Dell network. You should only allow traffic as needed from these ports.

- For logging: SYSLOG (UDP port 514) between the controller and syslog servers.
- For software upgrade or retrieving system logs: TFTP (UDP port 69) or FTP (TCP ports 21 and 22) between the controller and a software distribution server.
- If the controller is a PPTP VPN server, allow PPTP (UDP port 1723) and GRE (protocol 47) to the controller.
- If the controller is an L2TP VPN server, allow NAT-T (UDP port 4500), ISAKMP (UDP port 500) and ESP (protocol 50) to the controller.
- If a third-party network management system is used, allow SNMP (UDP ports 161 and 162) between the network management system and all controllers.
- For authentication with a RADIUS server: RADIUS (typically, UDP ports 1812 and 813, or 1645 and 1646) between the controller and the RADIUS server.
- For authentication with an LDAP server: LDAP (UDP port 389) or LDAPS (UDP port 636) between the controller and the LDAP server.

- For authentication with a TACACS+ server: TACACS (TCP port 49) between the controller and the TACACS+ server.
- For packet captures: UDP port 5555 from an AP to an Ethereal packet-capture station; UDP port 5000 from an AP to a Wildpackets packet-capture station.
- For telnet access: Telnet (TCP port 23) from the network administrator's computer to any AP, if “telnet enable” is present in the “ap location 0.0.0" section of the controller configuration.
- For External Services Interface (ESI): ICMP (protocol 1) and syslog (UDP port 514) between a controller and any ESI servers.
- For XML API: HTTP (TCP port 80) or HTTPS (TCP port 443) between a controller and an XML-API client.

Chapter 30
Remote Access Points

The Secure Remote Access Point Service allows AP users, at remote locations, to connect to a Dell controller over the Internet. Because the Internet is involved, data traffic between the controller and the remote AP is VPN encapsulated. That is, the traffic between the controller and AP is encrypted. Remote AP operations are supported on all of Dell's APs.

Topics in this chapter include:

- About Remote Access Points on page 629
- Configuring the Secure Remote Access Point Service on page 631
- Deploying a Branch Office/Home Office Solution on page 636
- Enabling Remote AP Advanced Configuration Options on page 641
- Understanding Split Tunneling on page 656
- Understanding Bridge on page 662
- Provisioning Wi-Fi Multimedia on page 667
- Reserving Uplink Bandwidth on page 667
- Provisioning 4G USB Modems on Remote Access Points on page 668
- Configuring W-IAP3WN and W-IAP3WNP Access Points on page 674
- Converting an IAP to RAP or CAP on page 674
- Enabling Bandwidth Contract Support for RAPs on page 676

About Remote Access Points

Remote APs connect to a controller using Extended Authentication and Internet Protocol Security (XAuth/IPSec). AP control and 802.11 data traffic are carried through this tunnel. Secure Remote Access Point Service extends the corporate office to the remote site. Remote users can use the same features as corporate office users. For example, voice over IP (VoIP) applications can be extended to remote sites while the servers and the PBX remain secure in the corporate office.

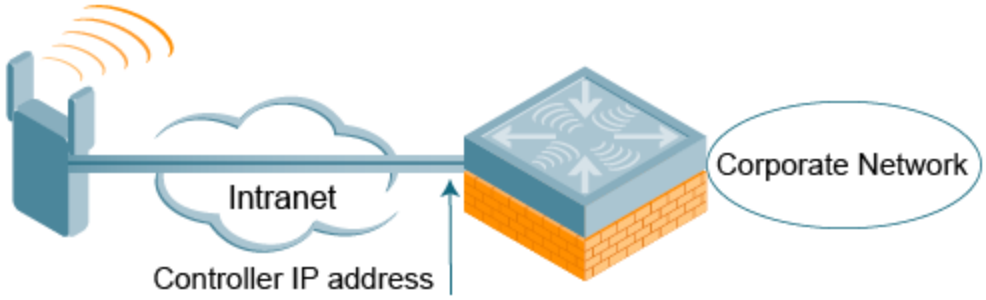
For both RAPs and CAPs, tunneled SSIDs will be brought down eight seconds after the AP detects that there is no connectivity to the controller. However, RAP bridge-mode SSIDs are configurable to stay up indefinitely (always-on / persistent). For CAP bridge-mode SSIDs, the CAP will be brought down after the keepalive times out (default 3.5 minutes).

Secure Remote Access Point Service can also be used to secure control traffic between an AP and the controller in a corporate environment. In this case, both the AP and controller are in the company's private address space.

The remote AP must be configured with the IPSec VPN tunnel termination point. Once the VPN tunnel is established, the AP bootstraps and becomes operational. The tunnel termination point used by the remote AP depends upon the AP deployment, as shown in the following scenarios:

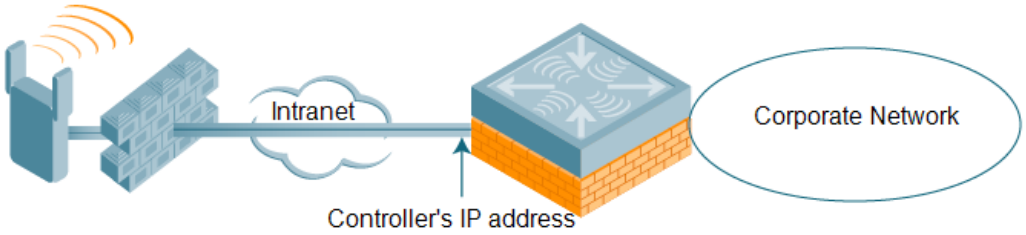
- Deployment Scenario 1: The remote AP and controller reside in a private network which secures AP-to-controller communication. (This deployment is recommended when AP-to-controller communications on a private network need to be secured.) In this scenario, the remote AP uses the controller's IP address on the private network to establish the IPSec VPN tunnel.

Figure 70 *Remote AP with a Private Network*



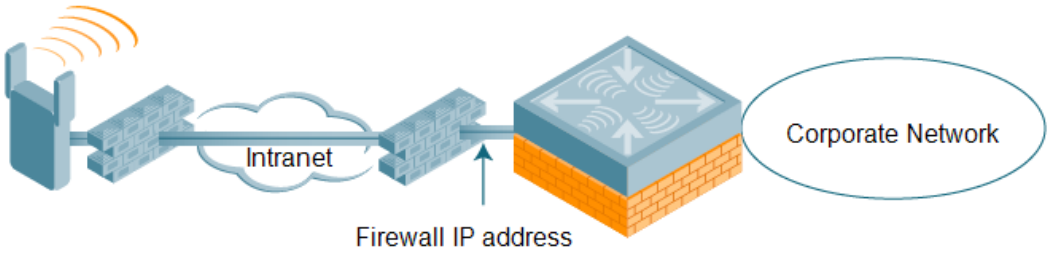
- Deployment Scenario 2: The remote AP is on the public network or behind a NAT device and the controller is on the public network. The remote AP must be configured with the tunnel termination point, which must be a publicly-routable IP address. In this scenario, a routable interface is configured on the controller in the DMZ. The remote AP uses the controller's IP address on the public network to establish the IPSec VPN tunnel.

Figure 71 *Remote AP with Controller on Public Network*



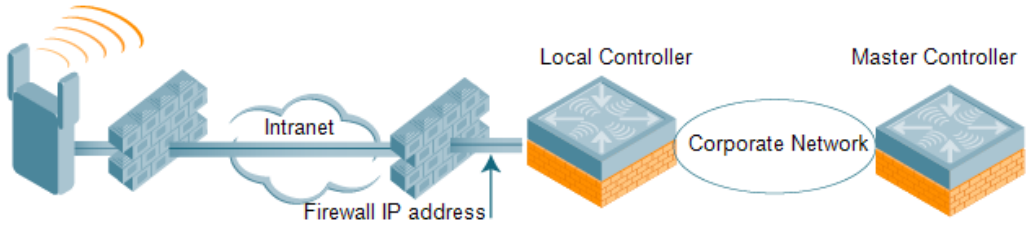
- Deployment Scenario 3: The remote AP is on the public network or behind a NAT device and the controller is also behind a NAT device. (This deployment is recommended for remote access.) The remote AP must be configured with the tunnel termination point, which must be a publicly-routable IP address. In this scenario, the remote AP uses the public IP address of the corporate firewall. The firewall forwards traffic to an existing interface on the controller. (The firewall must be configured to pass NAT-T traffic (UDP port 4500) to the controller.)

Figure 72 *Remote AP with Controller Behind Firewall*



In any of the described deployment scenarios, the IPSec VPN tunnel can be terminated on a local controller, with a master controller located elsewhere in the corporate network (Figure 73). The remote AP must be able to communicate with the master controller after the IPSec tunnel is established. Make sure that the L2TP IP pool configured on the local controller (from which the remote AP obtains its address) is reachable in the controller network by the master controller.

Figure 73 *Remote AP in a Multi-Controller Environment*



Configuring the Secure Remote Access Point Service

The tasks for configuring a Dell Access Points as a Secure Remote Access Point Service are:

- Configure a public IP address for the controller.

  You must install one or more AP licenses in the controller. There are several AP licenses available that support different maximum numbers of APs. The licenses are cumulative; each additional license installed increases the maximum number of APs supported by the controller.
- Configure the VPN server on the controller. The remote AP will be a VPN client to the server.
- Provision the AP with IPSec settings, including the username and password for the AP, before you install it at the remote location. You can also provision the RAP using the zero touch provisioning method. For more information, see Provisioning 4G USB Modems on Remote Access Points on page 668.

Configure a Public IP Address for the Controller

The remote AP requires an IP address to which it can connect to establish a VPN tunnel to the controller. This can be either a routable IP address you configure on the controller, or the address of an external router or firewall that forwards traffic to the controller. The following procedure describes how to create a DMZ address on the controller.

Using the WebUI to create a DMZ address

1. Navigate to the **Configuration** > **Network** > **VLANs** page.
2. Click **Add** to add a VLAN.
3. Enter the VLAN ID.
4. Select the port that belongs to this VLAN.
5. Click **Apply**.
6. Navigate to the **Configuration** > **Network** > **IP** page.
7. Click **Edit** for the VLAN you just created.
8. Enter the IP Address and Net Mask fields.
9. Click **Apply**.

Using CLI

Use the following commands:

```
vlan <id>
interface fastethernet <slot>/<port>
   switchport access vlan <id>
interface vlan <id>
   ip address <ipaddr> <mask>
```

Configure the NAT Device

Communication between the AP and the secure controller uses the UDP 4500 port. When both the controller and the AP are behind NAT devices, configure the AP to use the NAT device's public address as its master address. On the NAT device, you must enable NAT-T (UDP port 4500 only) and forward all packets to the public address of the NAT device on UDP port 4500 to the controller to ensure that the remote AP boots successfully.

Configure the VPN Server

This section describes how to configure the IPSec VPN server on the controller. For more details, see Virtual Private Networks on page 337. The remote AP will be a VPN client that connects to the VPN server on the controller.

Using the WebUI

1. Navigate to the **Configuration** > **Advanced Services** > **VPN Services** > **IPSec** page.
2. Select **Enable L2TP**.
3. Make sure that **PAP (Password Authentication Protocol)** is selected for Authentication Protocols.
4. To configure the L2TP IP pool, click **Add** in the **Address Pools** section. Configure the L2TP pool from which the APs will be assigned addresses, then click **Done**.



The size of the pool should correspond to the maximum number of APs that the controller is licensed to manage.

5. To configure an Internet Security Association and Key Management Protocol (ISAKMP) encrypted subnet and preshared key, click **Add** in the **IKE Shared Secrets** section and configure the preshared key. Click **Done** to return to the IPSec page.
6. Click **Apply**.

Using CLI

```
vpdn group l2tp
   ppp authentication PAP

ip local pool <pool> <start-ipaddr> <end-ipaddr>
crypto isakmp key <key> address <ipaddr> netmask <mask>
```

CHAP Authentication Support over PPPoE

RAPs can now establish a PPPoE session with a PPPoE server at the ISP side and get authenticated using the Challenge Handshake Authentication Protocol (CHAP). The PPPoE client running on a RAP is capable of handling the CHAP authentication requests from the PPPoE server.



The PPPoE client selects either the PAP or the CHAP credentials for the RAP authentication depending upon the request from the PPPoE server.

You can use the WebUI or the CLI to configure CHAP.

Using the WebUI to configure CHAP

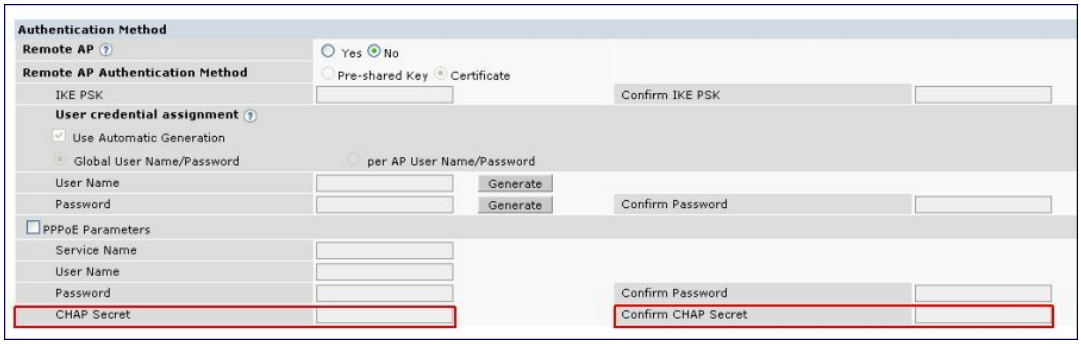
1. Navigate to the **Configuration** > **Wireless** > **AP Installation** page. The list of discovered APs are displayed on this page.
2. Select the AP you want to configure using CHAP and click **Provision** button.
3. Enter the **CHAP Secret** in the text box under **Authentication Method**.

NOTE

You can use all the special characters except question mark (?) and the space can be used within double quotes (" ").

4. Enter the **CHAP Secret** again in the **Confirm CHAP Secret** text box for confirmation.

Figure 74 *CHAP Authentication Using CHAP Secret*



5. Click **Apply and Reboot.**

Using the CLI to configure the CHAP

Use the following commands:

```
provision-ap pppoe-chap-secret <KEY>
   reprovision ap-name <name>
```

Configuring Certificate RAP

You can configure the remote AP to use the internal certificate for authentication. You can use the WebUI or CLI to configure the certificate RAP.

Using WebUI

1. Navigate to **Configuration** > **AP Installation** (under Wireless.)
2. Select the required remote AP under the **Provisioning** tab and then click **Provision**.
3. Select **Yes** for Remote AP and **Certificate** for Remote AP Authentication Method.
4. Click **Apply and Reboot** to apply the configuration and reboot the AP as certificate RAP.

Using CLI

Use the following command:

```
   local-userdb-ap whitelist-db rap add <mac-address>
```

Creating a Remote AP Whitelist

If you use the Zero Touch provisioning method to provision the certificate RAP, then you must create a remote AP whitelist. For more information on Zero Touch Provisioning of the RAP, see Provisioning 4G USB Modems on Remote Access Points on page 668.

Remote AP whitelist is the list of approved APs that can be provisioned on your controller. To create a remote AP whitelist:

1. Navigate to **Configuration** > **AP Installation** (under Wireless) and then click the **RAP Whitelist** tab on the right side.
2. Click **New** and provide the following details:
   - **AP MAC Address**—mandatory parameter. Enter the MAC address of the AP.

- **Username**—enter a username that is used when the AP is provisioned.
- **AP Group**—select a group to add the AP.
- **AP Name**—enter a name for the AP. If you do not enter an AP name, the MAC address will be used instead.
- **Description**—enter a text description for the AP
- **IP-Address**—enter an IP address for the AP.

3. Click **Add** to add the remote AP to the whitelist.

Configuring PSK RAP

You can use Pre-Shared Key (PSK) authentication to provision an individual remote AP or a group of remote APs using an Internet Key Exchange Pre-Shared Key (IKE PSK).

To configure the PSK RAP using the WebUI:

1. Navigate to the **Configuration** > **Wireless** > **AP Installation** > **Provisioning** window.
2. Click the checkbox by the AP you want to provision, then click **Provision**. The Provisioning window opens.
3. Select **Yes** for the **Remote AP** option
4. In the **Remote IP Authentication Method** section, select **Pre-shared key**.
5. Enter and confirm the pre-shared key (IKE PSK).
6. In the **User credential assignment** section, specify if you want to use a **Global User Name/password** or a **Per AP User Name/Password**.
   a. If you use the **Per AP User Names/Passwords** option, each RAP is given its own username and password.
   b. If you use the **Global User Name/Password** option, all selected RAPs are given the same (shared) username and password.
7. Enter the user name, and enter and confirm the password. If you want the controller to automatically generate a user name and password, select **Use Automatic Generation**, then click **Generate by the User Name and Password** fields.

Add the user to the internal database

You can add the user to the internal database using the WebUI or CLI.

Using WebUI

1. Navigate to the **Configuration** > **Security** > **Authentication** > **Servers** page.
2. Select **Internal DB**.
3. Click **Add User** in the Users section. The user configuration page displays.
4. Enter the username and password.
5. Click **Enabled** to activate this entry on creation.
6. Click **Apply** . Note that the configuration does not take effect until you perform this step.
7. At the **Servers** page, click **Apply**.

Using CLI

Use the following command:

```
local-userdb add username rapuser1 password <password>
```

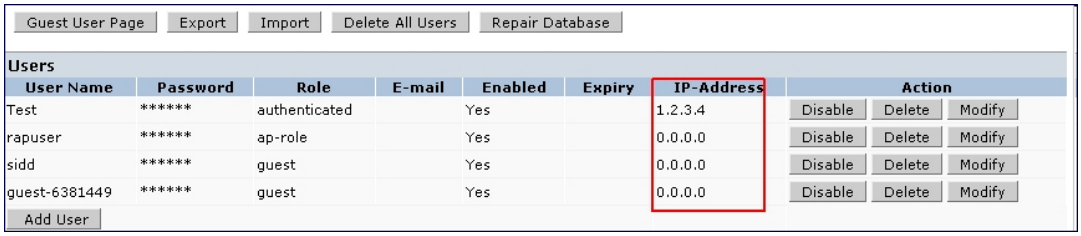
RAP Static Inner IP Address

The RAP static inner IP address feature assigns a static inner IP address to a remote access point (RAP). A new *remote-IP address* parameter is added to the existing configuration commands.

Using the WebUI

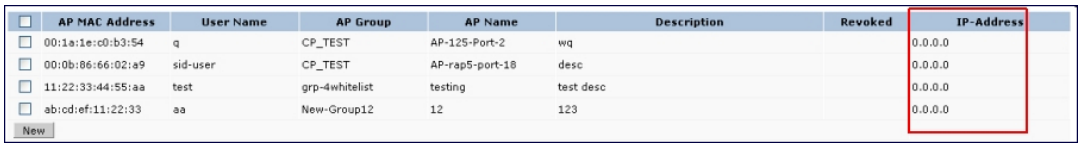
To view IP address parameter in the local database, navigate to the **Configuration** > **Security** > **Authentication** > **Servers** > **Internal DB** page.

Figure 75 *IP-Address parameter in the local database*



To view IP-address parameter in the RAP Whitelist, navigate to the **Wireless** > **AP Installation** > **RAP Whitelist** page.

Figure 76 *IP-Address parameter in the RAP Whitelist*



Using the CLI

```
local-userdb add {generate-username|username <name>} {generate-password|password
<password>} {remote-ip <remote-ip>}
local-userdb modify {username < name>} {remote-ip <remote-ip>}
local-userdb-ap whitelist-db rap add {mac-address <address>}{ap-group <ap_group>}{remote-ip
<remote-ip>}
local-userdb-ap whitelist-db rap modify {mac-address <address>} {remote-ip<remote-ip>}
```



You cannot configure the IP-Address parameter using the WebUI.

Provision the AP

You need to configure the VPN client settings on the AP to instruct the AP to use IPSec to connect to the controller. You can provision the remote AP and give it to users and allow remote users to provision AP at their home. This method of provisioning is referred as Zero Touch Provisioning. See Provisioning 4G USB Modems on Remote Access Points on page 668 for more information about Zero Touch Provisioning of remote AP.

You must provision the AP before you install it at its remote location. To provision the AP, the AP must be physically connected to the local network or directly connected to the controller. When connected and powered on, the AP must also be able to obtain an IP address from a DHCP server on the local network or from the controller.

If your configuration has an internal LMS IP address, remote APs may attempt to switch over to the LMS IP address, which is not reachable from the Internet. For remote APs, ensure that the LMS IP address in the AP system profile for the AP group has an externally routable IP address.

Reprovisioning the AP causes it to automatically reboot. The easiest way to provision an AP is to use the Provisioning page in the WebUI, as described in the following steps:

1. Navigate to the **Configuration** > **Wireless** > **AP Installation** > **Provisioning** page. Select the remote AP and click **Provision**.
2. Under **Authentication Method**, select **IPSec Parameters**. Enter the **Internet Key Exchange (IKE) Pre-Shared Key (PSK)**, **username**, and **password**.



The username and password you enter must match the username and password configured on the authentication server for the remote AP.

3. Under Master Discovery, set the Master IP Address as shown below:

Deployment Scenario	Master IP Address Value
Deployment 1	Controller IP address
Deployment 2	Controller public IP address
Deployment 3	Public address of the NAT device to which the controller is connected



The username and password you enter must match the username and password configured on the authentication server for the remote AP.

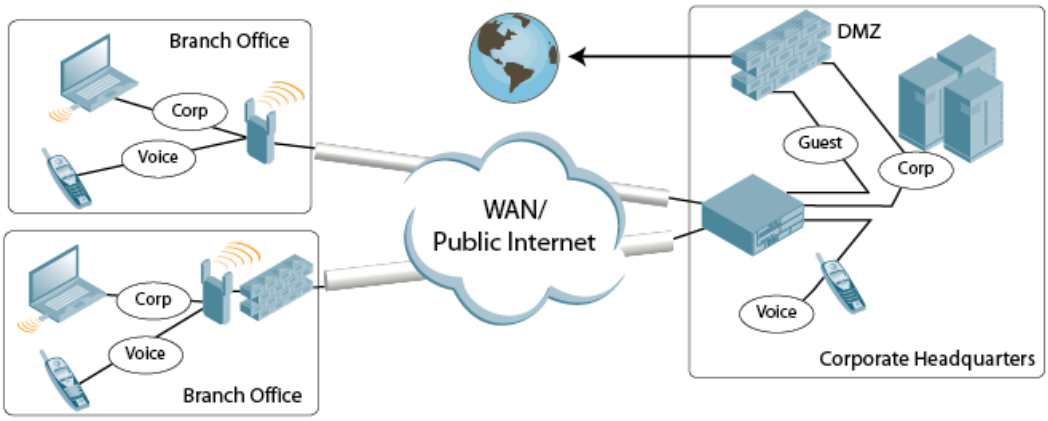
4. Under **IP Settings**, make sure that **Obtain IP Address Using DHCP** is selected.
5. Click **Apply and Reboot**.

Deploying a Branch Office/Home Office Solution

In a branch office, the AP is deployed in a separate IP network from the corporate network. Typically, there are one or two NAT devices between the two networks. Branch office users need access to corporate resources such as printers and servers, but traffic to and from these resources must not impact the corporate head office.

Figure 77 is a graphic representation of a remote AP in a branch or home office, with a single controller providing access to both a corporate WLAN and a branch office WLAN.

Figure 77 *Remote AP with Single Controller*



Branch office users want continued operation of the branch office WLAN, even if the link to the corporate network goes down. The branch office AP solves these requirements by providing the following capabilities on the branch office WLAN:

- Local termination of 802.11 management frames which provides survivability of the branch office WLAN.
- All 802.1x authenticator functionality is implemented in the AP. The controller is used as a RADIUS pass-through when the authenticator has to communicate with a RADIUS server (which also supports survivability).
- 802.11 encryption/decryption is in the AP to provide access to local resources.
- Local bridging of client traffic connected to the WLAN or to an AP 70 enet1 port to provide access to local resources.

Provisioning the Branch Office AP

You can provision the remote AP either using the controller or using the Zero Touch Provisioning method. For more information on controller provisioning, see Configuring Installed APs on page 500. For more information on Zero Touch Provisioning, see Provisioning 4G USB Modems on Remote Access Points on page 668.

Configuring the Branch Office AP

- Specify forward mode for the Extended Service Set Identifier (ESSID) in the virtual AP profile
- Specify remote AP operation in the virtual AP profile (The remote AP operates in standard mode by default.)
- Set how long the AP stays up after connectivity to controller has gone down in the SSID profile
- Set the VLAN ID in the virtual AP profile
- Set the native VLAN ID in the AP system profile
- Set forward mode for enet1 port



Remote APs support 802.1q VLAN tagging. Data from the remote AP will be tagged on the wired side.

Troubleshooting Remote AP

The following WebUI options are available to troubleshoot issues with remote AP:

- Using local debugging feature
- Viewing the remote AP summary report
- Viewing remote AP connectivity report

- Using remote AP diagnostic options

Local Debugging

Local debugging is a WebUI feature that allows end users to perform diagnostics and view the status of their remote AP through a wired or wireless client. This feature is useful for troubleshooting connectivity problems on remote APs and performing throughput tests. There are three tabs in the **Local Debugging** WebUI window; **Summary**, **Connectivity**, and **Diagnostics**. Each tab displays different information for the AP, but all three tabs include a **Generate & save support file** link that, when clicked, will automatically generate a **support.tgz** file that can be sent to a corporate IT department for additional analysis and debugging.



Starting from Dell Networking W-Series ArubaOS 6.4.x, snapshot of the bridge, acl, session, user, and arp tables, current processes, memory, and kernel debug messages are captured in a single **rap_debug.txt** file which is bundled along with **support.tgz** file.

Remote AP Summary

The **Summary** tab has two views; basic and advanced. Click the **basic** or **advanced** links at the top of this tab to toggle between the two views. The table below shows the information displayed for both the basic and advanced views of the **Summary** tab.

Table 121: *RAP Console Summary Tab Information*

Summary Table Name	Basic View Information	Advanced View Information
Wired Ports Status	• **Port**: port numbers of the wired ports on the AP • **Status**: current status of each port (*Connected*, *LinkDown* or *Disabled*).	The advanced view of the Wired Access Ports table displays the following data: • **Port**: port numbers of the wired ports on the AP • **Status**: current status of each port (*Connected*, *LinkDown* or *Disabled*) • **MAC Address:** MAC address of the wired port • **Speed**: speed of the link • **Duplex Type**: duplex mode of the link, full or half • **Forwarding mode**: forwarding mode for the port: *Bridge*, *Tunnel* or *Split Tunnel* • **Users**: fumber of users accessing each port • **Rx Packets**: number of packets received on the port • **Tx packets**: number of packets transmitted via the port
Wireless SSIDs	• **SSID**: Name of the SSID. • **Status**: SSID Status (up, down, or disabled). • **Band**: Radio band available on the SSID.	• **SSID**: name of the SSID • **Status**: SSID Status (up, down, or disabled). • **Band**: radio band available on the SSID • **Channel**: channel used on the radio band • **BSSID**: BSSID of the wireless SSID • **Forwarding Mode**: forwarding mode used by the Wireless SSID (Bridge, Tunnel or Split-Tunnel) • **EIRP**: equivalent Isotropic Radiated Power, in dBm • **Noise floor**: residual background noise detected by an AP. Noise seen by an AP is reported as -dBm Therefore, a noise floor of -100 dBm is smaller (lower) than a noise floor of -50 dBm. • **Users**: number of users on the radio band • **Rx Packets**: number of packets received on the BSSID • **Tx packets**: number of packets transmitted via the BSSID
Wired Users	• **MAC Address**: MAC address of the wired user. • **IP address**: IP address of the wired user.	• **MAC Address**: MAC address of the wired user. • **IP address**: IP address of the wired user. • **Port**: AP port used by the wired user.
Wireless User	• **MAC Address**: MAC address of the wireless user. • **IP address**: IP address of the wireless user.	• **MAC Address**: MAC address of the wired user • **IP address**: IP address of the wired user • **SSID**: name of the SSID • **BSSID**: BSSID of the wireless user • **Assoc State**: shows if the user is

Summary Table Name	Basic View Information	Advanced View Information
		associated or just authorized • **Auth**: Type of authentication: WPA, 802.1x, none, open, or shared • **Encryption**: encryption type used by the wireless user • **Band**: radio band used by the wireless client • **RSSI**: Receive Signal Strength Indicator (RSSI) value displayed in the output of this command represents signal strength as a signal to noise ratio.
Device Info	• **Type**: AP device/model type. • **Name**: Name assigned to the AP. • **Wired MAC address**: MAC address of the wired port. • **Serial #**: AP serial number. • **Tunnel IP address**: IP address of the tunnel between the AP and controller. • **Software Version**: Software version currently running on the AP. • **Uptime**: Amount of time the AP has been active since it was last reset. • **Master**: IP address of the master controller. • **lms**: IP address of the local controller.	N/A
Uplink Info	The Uplink Info table can display some or all of the following information for your remote AP, depending upon whether a link is active and the number of links supported by the AP. Active uplink information, including: • Interface name • Port speed • IP address Standby link information, including: • Name (3G) • **Device connected** (yes/no) • **Provisioned** (yes/no) • IP address • Device • User • Password	N/A

Multihoming on remote AP (RAP)

You can uplink a RAP as an Ethernet or a USB based modem. These uplinks can be used as a backup link if the primary link fails. The uplink becomes active based on the order of priority configured on the RAP. The RAP switches back to the primary link when the primary connection is restored.

For information on provisioning the RAP using the USB based modem, see Provisioning 4G USB Modems on Remote Access Points on page 668.

Seamless failover from backup link to primary link on RAP

RAPs can failover from a backup link to a primary link without much disruption to traffic. Also the failover is performed only if the controller is reachable via the primary link.

Remote AP Connectivity

The information shown on the **Connectivity** tab will vary, depending upon the current status of the remote AP. If a remote AP has been successfully provisioned and connected, it should display some or all of the information in Table 122.

Table 122: *RAP Console Connectivity Tab Information*

Data	Description
Uplink status	Shows if the link connected failed. If the link is connected, the Uplink status also displays the name of the interface.
IP Information	If the AP has successfully received an IP address, this data row will show the AP's IP address, subnet mask, and gateway IP address.
Gateway Connectivity	If successful, this item also shows the percentage of packet loss for data received from the gateway.
TPM Certificates	If successful, the AP has a Trusted Platform Module (TPM) certificate.
Master Connectivity	Shows if the AP was able to connect to the master controller. This item also shows the IP address to which the AP attempted to connect, and, if the AP did connect successfully, the link used to connect to that controller.
LMS Connectivity	Shows if the AP was able to connect to a local controller. This item also shows the IP address to which the AP attempted to connect, and, if the AP did connect successfully, the link used to connect to that controller.

The top of the **Connectivity** tab has a **Refresh** link that allows users to refresh the data on their screen. Additional information at the bottom of this tab shows the date, time, and reason the remote AP last rebooted. The **Reboot RAP Now** button reboots the remote AP.

Remote AP Diagnostics

Use the **Diagnostics** tab to view log files, or run diagnostic tests that can help the IT department troubleshoot errors. Use the **Reboot AP Now** button at the bottom of the Diagnostic window to reboot the remote AP.

To run a diagnostic test on a remote AP:

1. Access the RAP console, and click the **Diagnostics** tab.
2. Click the **Test** drop-down list and select **Ping**, **Traceroute, NSLookup**, or **Throughput**.

   The *ping* and *traceroute* tests require that you enter a network destination in the form of an IP address or fully-qualified domain name, and select either **bridge** or **tunnel** mode for the test. *The NSLookup* diagnostic test requires that you enter a destination only. The *throughput* test checks the throughput of the link between the AP and the controller, and does not require any additional test configuration settings.
3. Click **OK** to start the test. The results of the test will appear in the **Diagnostics** window.

To display log files in a separate browser window, click the **logs** drop-down list at the upper right corner of the **Diagnostics** window, and select any of the log file name. The type of log files available will vary, depending upon your remote AP configuration.

Enabling Remote AP Advanced Configuration Options

This section describes the following features designed to enhance your remote AP configuration:

- Understanding Remote AP Modes of Operation on page 642
- Working in Fallback Mode on page 644

- Specifying the DNS Controller Setting on page 652
- Backup Controller List on page 653
- Configuring Remote AP Failback on page 654
- Working with Access Control Lists and Firewall Policies on page 656
- Understanding Split Tunneling on page 656
- Provisioning Wi-Fi Multimedia on page 667



The information in this section assumes you have already configured the remote AP functionality, as described in Configuring the Secure Remote Access Point Service on page 631.

Understanding Remote AP Modes of Operation

Table 123 summarizes the different remote AP modes of operation. You specify both the forward mode setting (which controls whether 802.11 frames are tunneled to the controller using GRE, bridged to the local Ethernet LAN, or a combination thereof) and the remote AP mode of operation (when the virtual AP operates on a remote AP) in the virtual AP profile.

The column on the left of the table lists the remote AP operation settings. The row across the top of the table lists the forward mode settings. To understand how these settings work in concert, scan the desired remote AP operation with the forward mode setting, and read the information in the appropriate table cell.

The “all” column and row lists features that all remote AP operation and forward mode settings have in common regardless of other settings. For example, at the intersection of “all” and “bridge,” the description outlines what happens in bridge mode regardless of the remote AP mode of operation.

Table 123: *Remote AP Modes of Operation and Behavior*

Remote AP Operation Setting	Forward Mode Setting				
	all	**bridge**	**split-tunnel**	**tunnel**	decrypt-tunnel
all		Management frames on the AP. Frames are bridged between wired and wireless interfaces. No frames are tunneled to the controller. Station acquires its IP address locally from an external DHCP server.	Management frames on the AP. Frames are either GRE tunneled to the controller to a trusted tunnel or NATed and bridged on the wired interface according to user role and session ACL. Typically, the station obtains an IP address from a VLAN on the controller. Typically, the AP has ACLs that forward corporate traffic through the tunnel and source NAT the non-corporate traffic to the Internet.	Frames are GRE tunneled to the controller to an untrusted tunnel. 100% of station frames are tunneled to the controller.	Management frames on the AP. Frames are always GRE tunneled to controller.
always	ESSID is always up when the AP is up regardless of whether the controller is reachable. Supports PSK ESSID only. SSID configuration stored in flash on AP.	Provides an SSID that is always available for local access.	Not supported	Not supported	Not supported
	all	**bridge**	**split-tunnel**	**tunnel**	

Remote AP Operation Setting	Forward Mode Setting				
backup	ESSID is only up when the controller is unreachable. Supports PSK ESSID only. SSID configuration stored in flash on AP.	Provides a backup SSID for local access only when the controller is unreachable.	Not supported	Not supported	Not supported
persistent	ESSID is up when the AP contacts the controller and stays up if connectivity is disrupted with the controller. SSID configuration obtained from the controller. Designed for 802.1x SSIDs.	Same behavior as standard, described below, except the ESSID is up if connectivity to the controller is lost.	Not supported	Not supported	Not supported
standard	ESSID is up only when there is connectivity with the controller. SSID configuration obtained from the controller.	Behaves like a classic Dell branch office AP. Provides a bridged ESSID that is configured from the controller and stays up if there is controller connectivity.	Split tunneling mode	Classic Dell thin AP operation	Decrypt tunnel mode

Working in Fallback Mode

The fallback mode (also known as backup configuration) operates the remote AP if the master controller or the configured primary and backup LMS are unreachable. The remote AP saves configuration information that allows it to operate autonomously using one or more SSIDs in local bridging mode, while supporting open association or encryption with PSKs. You can also use the backup configuration if you experience network connectivity issues, such as the WAN link or the central data center becoming unavailable. With the backup configuration, the remote site does not go down if the WAN link fails or the data center is unavailable.

You define the backup configuration in the virtual AP profile on the controller. The remote AP checks for configuration updates each time it establishes a connection with the controller. If the remote AP detects a change, it downloads the configuration changes.

The following remote AP backup configuration options define when the SSID is advertised (refer to Table 123 for more information):

- Always—Permanently enables the virtual AP. Recommended for bridge SSIDs.
- Backup—Enables the virtual AP if the remote AP cannot connect to the controller. This SSID is advertised until the controller is reachable. Recommended for bridge SSIDs.

- Persistent—Permanently enables the virtual AP after the remote AP initially connects to the controller. Recommended for 802.1x SSIDs.
- Standard—Enables the virtual AP when the remote AP connects to the controller. Recommended for 802.1x, tunneled, and split-tunneled SSIDs. This is the default behavior.

While using the backup configuration, the remote AP periodically retries its IPSec tunnel to the controller. If you configure the remote AP in backup mode, and a connection to the controller is re-established, the remote AP stops using the backup configuration and immediately brings up the standard remote AP configuration. If you configure the remote AP in always or persistent mode, the backup configuration remains active after the IPSec tunnel to the controller has been re-established.

Backup Configuration Behavior for Wired Ports

If the connection between the remote AP and the controller is disconnected, the remote AP will be exhibit the following behavior:

- All access ports on the remote AP will be moved to bridge forwarding mode ,irrespective of their original forwarding mode..
- Clients will receive an IP address from the remote AP's DHCP server.
- Clients will have complete access to Remote AP's uplink network. You cannot enforce or modify any access control policies on the clients connected in this mode.

This section describes the following topics:

- Configuring Fallback Mode on page 645
- Configuring the DHCP Server on the Remote AP on page 647
- Configuring Advanced Backup Options on page 649

Configuring Fallback Mode

To configure the fallback mode, you must:

- Configure the AAA profile
- Configure the virtual AP profile

Configuring the AAA Profile for Fallback Mode in the WebUI

The AAA profile defines the authentication method and the default user role for unauthenticated users:

1. Navigate to the **Security > Authentication > AAA Profiles** page. From the AAA Profiles Summary list, click **Add**.
2. Enter the AAA profile name, then click **Add**.
3. Select the AAA profile that you just created:
   a. For **Initial role**, select the appropriate role (for example, “logon”).
   b. For **802.1X Authentication Default Role**, select the appropriate role (for example, “default”), then click **Apply**.
   c. Under the AAA profile that you created, locate **802.1x Authentication Server Group**, and select the authentication server group to use (for example “default”), then click **Apply**.



If you need to create an 802.1x authentication server group, select **new** from the 802.1X Authentication Server Group drop-down list, and enter the appropriate parameters.

   d. Under the AAA profile that you created, locate 802.1X Authentication Profile, and select the profile to use (for example, “default”), then click **Apply**.

NOTE

If you need to create an 802.1x authentication profile, select new from the **802.1X Authentication Profile** drop-down list, and enter the appropriate parameters.

Configuring the AAA Profile for Fallback Mode in the CLI

Use the following command:

```
aaa profile <name>
   initial-role <role>
   authentication-dot1x <dot1x-profile>
   dot1x-default-role <role>
   dot1x-server-group <group>
```

Configuring the Virtual AP Profile for Fallback Mode in the WebUI

To configure virtual AP profile:

- Set the remote AP operation to **always**, **backup**, or **persistent**.
- Create and apply the applicable SSID profile.

  The SSID profile for the backup configuration in always, backup, or persistent mode must be a bridge SSID. When configuring the virtual AP profile, specify forward mode as **bridge**.

  The SSID profile for the backup configuration in standard mode can be a bridge, tunnel, or split tunnel SSID. When configuring the virtual AP profile, specify forward mode as **bridge**, **tunnel**, or **split tunnel**.

NOTE

When creating a new virtual AP profile In the WebUI, you can also configure the SSID at the same time. For information about AP profiles, see Understanding AP Configuration Profiles on page 488.

1. Navigate to the **Configuration > Wireless > AP Configuration** page. Select either the **AP Group** or **AP Specific** tab. Click **Edit** for the AP group or AP name.
2. Under Profiles, select **Wireless LAN**, then **Virtual AP**.
3. To create a new virtual AP profile in the WebUI, select **New** from the **Add a profile** drop-down menu. Enter the name for the virtual AP profile, and click **Add**.

NOTE

Whenever you create a new virtual AP profile in the WebUI, the profile automatically contains the “default” SSID profile with the default ESSID. You must configure a new ESSID and SSID profile for the virtual AP profile before you apply the profile.

   a. In the Profile Details entry for the new virtual AP profile, go to the **AAA Profile** drop-down list and select the previously configured AAA profile (for example, **logon**). The AAA Profile pop-up window appears.
   b. To set the AAA profile and close the pop-up window, Click **Apply.**
   c. In the Profile Details entry for the new virtual AP profile, select **NEW** from the **SSID Profile** drop-down menu. The SSID Profile pop-up window displays to allow you to configure the SSID profile.
   d. Enter the name for the SSID profile (for example, **backup**).
   e. Under Network, enter a name in the Network Name (SSID) field (for example, **backup-psk**).
   f. Under Security, select the network authentication and encryption methods (for example, wpa-psk-tkip, with the passphrase **remote123**).
   g. To set the SSID profile and close the pop-up window, click **Apply**.
4. At the bottom of the **Profile Details** window, Click **Apply**.
5. Click the new virtual AP name in the **Profiles** list or the **Profile Details** to display configuration parameters.
6. Under Profile Details, do the following:
   a. Make sure **Virtual AP enable** is selected.

b. From the **VLAN** drop-down menu, select the VLAN ID to use for the virtual AP profile.

c. From the **Forward mode** drop-down menu, select **bridge**.

d. From the **Remote-AP Operation** drop-down menu, select **always**, **backup**, or **persistent**. The default is standard. Click **Apply**.

Configuring the Virtual AP Profile for Fallback Mode in the CLI

```
wlan ssid-profile <profile>
   essid <name>
   opmode <method>
   wpa-passphrase <string> (if necessary)

wlan virtual-ap <name>
   ssid-profile <profile>
   vlan <vlan>
   forward-mode bridge
   aaa-profile <name>
   rap-operation {always|backup|persistent}

ap-group <name>
   virtual-ap <name>
```

or

```
ap-name <name>
   virtual-ap <name>
```

Configuring the DHCP Server on the Remote AP

You can configure the internal DHCP server on the remote AP to provide an IP address for the backup SSID if the controller is unreachable. If configured, the remote AP DHCP server intercepts all DHCP requests and assigns an IP address from the configured DHCP pool.

To configure the remote AP DHCP server:

- Enter the VLAN ID for the remote AP DHCP VLAN in the AP system profile. This VLAN enables the DHCP server on the AP (also known as the remote AP DHCP server VLAN). If you enter the native VLAN ID, the DHCP server is not configured and is unavailable.
- Specify the DHCP IP address pool and netmask. The AP assigns IP addresses from the DHCP pool 192.168.11.0/24 by default, with an IP address range from 192.168.11.2 through 192.168.11.254. You can manually define the DHCP IP address pool and netmask based on your network design and IP address scheme.
- Specify the IP address of the DHCP server, DHCP router, and the DHCP DNS server. The AP uses IP address 192.168.11.1 for the DHCP server, the DHCP router, and the DHCP DNS server by default.
- Enter the amount of days the assigned IP address is valid (also known as the remote AP DHCP lease). The lease does not expire by default, which means the IP address is always valid.
- Assign the VLAN ID for the remote AP DHCP VLAN to a virtual AP profile. When a client connects to that virtual AP profile, the AP assigns the IP address from the DHCP pool.



The following is a high-level description of the steps required to configure the DHCP server on the remote AP. The steps assume you have already created the virtual AP profile, AAA profile, SSID profile, and other settings for your remote AP operation (for information about the backup configuration, see Configuring Fallback Mode on page 645).

Using the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. Select either the **AP Group** or **AP Specific** tab. Click **Edit** for the AP group or AP name.

3. Under **Profiles**, select **AP** to display the AP profiles.
4. Select the AP system profile you want to modify.
5. Under Profile Details:
   a. At the **LMS IP** field, enter the LMS IP address.
   b. At the **Master controller IP address** field, enter the master controller IP address.
   c. At the **Remote-AP DHCP Server VLAN** field, enter the VLAN ID of the backup configuration virtual AP VLAN.
   d. At the **Remote-AP DHCP Server ID** field, enter the IP address for the DHCP server.
   e. At the **Remote-AP DHCP Default Router** field, enter the IP address for the default DHCP router.
   f. At the **Remote-AP DHCP DNS Server** list, enter an IP address in the field to right and click **Add**. You can add multiple IP addresses the same way. To delete an IP address, select an IP address from the list and click **Delete**.
   g. Specify the DHCP IP address pool. This configures the pool of IP addresses from which the remote AP uses to assign IP addresses.
      —At the **Remote-AP DHCP Pool Start** field, enter the first IP address of the pool.
      —At the **Remote-AP-DHCP Pool End** field, enter the last IP address of the pool.
      —At the **Remote-AP-DHCP Pool Netmask** field, enter the netmask.
   h. At the **Remote-AP DHCP Lease Time** field, specify the amount of time the IP address is valid.
6. Click **Apply**.
7. Under **Profiles**, select **Wireless LAN**, then **Virtual AP**, then the virtual AP profile you want to configure.
8. Under **Profile Details**, at the VLAN drop-list, select the VLAN ID of the remote AP DHCP VLAN, click the left arrow to move the VLAN ID to the VLAN field, and click **Apply**.

Using CLI

Use the following commands:

```
ap system-profile <name>
   lms-ip <ipaddr>
   master-ip <ipaddr>
   rap-dhcp-default-router <ipaddr>
   rap-dhcp-dns-server <ipaddr>
   rap-dhcp-lease <days>
   rap-dhcp-pool-end <ipaddr>
   rap-dhacp-pool-netmask <netmask>
   rap-dhcp-pool-start <ipaddr>
   rap-dhcp-server-id <ipaddr>
   rap-dhcp-server-vlan <vlan>

wlan virtual-ap <name>
   ssid-profile <profile>
   vlan <vlan>
   forward-mode bridge
   aaa-profile <name>
   rap-operation {always|backup|persistent}

ap-group <name>
   ap-system-profile <name>
   virtual-ap <name>
```

or

```
ap-name <name>
   ap-system-profile <name>
   virtual-ap <name>
```

Configuring Advanced Backup Options

You can also use the backup configuration (fallback mode) to allow the remote AP to pass through a captive portal, such as network access in a hotel, airport, or other public network, to access the corporate network. For this scenario:

- Define a session ACL for the bridge SSID to source NAT all user traffic, except DHCP. For example, use **any any svc-dhcp permit** followed by **any any any route src-nat**. Apply the session ACL to a remote AP user role.
- Configure the AAA profile. Make sure the initial role contains the session ACL previously configured.

  The AAA profile defines the authentication method and the default user role.



802.1X and PSK authentication is supported when configuring bridge or split tunnel modes.

- Configure the virtual AP profile for the backup configuration:
  - Set the remote AP operation to **always** or **backup**.
  - Create and apply the applicable SSID profile.
  - Configure a bridge SSID for the backup configuration. In the virtual AP profile, specify forward mode as **bridge**.

  For more information about the backup configuration, see Configuring Fallback Mode on page 645.
- Enter the remote AP DHCP server parameters in the AP system profile. For more information about the parameters, see Configuring the DHCP Server on the Remote AP on page 647.

  If you use a local DHCP server to obtain IP addresses, you must define one additional ACL to permit traffic between clients without source NATing the traffic. Using the previously configured ACL, add **user alias internal-network any permit** before **any any any route src-nat**.
- Connect the remote AP to the available public network (for example, a hotel or airport network).

  The remote AP advertises the backup SSID so the wireless client can connect and obtain an IP address from the available DHCP server.



The client can obtain an IP address from the public network, for example a hotel or airport, or from the DHCP server on the remote AP.

  After obtaining an IP address, the wireless client can connect and access the corporate network and bring up the configured corporate SSIDs.

The following is a high-level description of what is needed to configure the remote AP to pass through a captive portal and access the corporate controller. This information assumes you are familiar with configuring session ACLs, AAA profiles, virtual APs, and AP system profiles and highlights the modified parameters.

Configuring the Session ACL in the WebUI

1. Navigate to the **Configuration** > **Security** > **Access Control** > **Policies** page.
2. Click **Add** to create a new policy.
3. Enter the policy name in the **Policy Name** field.
4. From the **Policy Type** drop-down list, select **IPv4 Session**.
5. To create the first rule:
   a. Under **Rules**, click **Add**.
   b. Under **Source**, select **any**.
   c. Under **Destination**, select **any**.
   d. Under **Service**, select **service**. In the service drop-down list, select **svc-dhcp**.

e. Under **Action**, select **permit**.
f. Click **Add**.

6. To create the next rule:
   a. Under **Rules**, click **Add**.
   b. Under **Source**, select **any**.
   c. Under **Destination**, select **any**.
   d. Under **Service**, select **any**.
   e. Under **Action**, select **route**, and select the **src-nat** checkbox.
   f. Click **Add**.
7. Click **Apply**.



If you use a local DHCP server to obtain IP addresses, you must define one additional ACL to permit traffic between clients without source NATing the traffic. Add `user alias internal-network any permit` before `any any any route src-nat`.

8. Click the **User Roles** tab.
   a. Click **Add**.
   b. Enter the Role Name.
   c. Click **Add** under Firewall Policies.
   d. In the **Choose from Configured Policies** menu, select the policy you just created.
   e. Click **Done**.

Configuring the AAA Profile in the WebUI

1. Navigate to the **Security** > **Authentication** > **AAA Profiles** page. From the AAA **Profiles Summary** list, click **Add**.
2. Enter the AAA profile name, then click **Add**.
3. Select the AAA profile that you just created:
   a. For Initial role, select the user role you just created.
   b. For 802.1X Authentication Default Role, select the appropriate role for your remote AP configuration, then click **Apply**.
   c. Under the AAA profile that you created, locate **802.1x Authentication Server Group**, and select the authentication server group to use for your remote AP configuration, then click **Apply**.



If you need to create an 802.1x authentication server group, select **new** from the **802.1X Authentication Server Group** drop-down list, and enter the appropriate parameters.

   d. Under the AAA profile that you created, locate **802.1X Authentication Profile**, select the profile to use for your remote AP configuration, then click **Apply**.

Defining the Backup Configuration in the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration** page. Select either the **AP Group** or **AP Specific** tab. Click **Edit** for the AP group or AP name.
2. Under **Profiles**, select **Wireless LAN**, then **Virtual AP**.
3. To create a new virtual AP profile in the WebUI, select **New** from the **Add a profile** drop-down menu. Enter the name for the virtual AP profile, and click **Add**.

NOTE

Whenever you create a new virtual AP profile in the WebUI, the profile automatically contains the "default" SSID profile with the default ESSID. You must configure a new ESSID and SSID profile for the virtual AP profile before you apply the profile.

a. In the **Profile Details** entry for the new virtual AP profile, go to the **AAA Profile** drop-down list and select the previously configured AAA profile. The AAA Profile pop-up window appears.
b. To set the AAA profile and close the pop-up window, Click **Apply.**
c. In the **Profile Details** entry for the new virtual AP profile, select **NEW** from the **SSID Profile** drop-down menu. The SSID Profile pop-up window displays to allow you to configure the SSID profile.
d. Enter the name for the SSID profile.
e. Under **Network**, enter a name in the Network Name (SSID) field.
f. Under **Security**, select the network authentication and encryption methods.
g. To set the SSID profile and close the pop-up window, click **Apply**.

4. At the bottom of the Profile Details window, Click **Apply**.
5. Click the new virtual AP name in the Profiles list or the Profile Details to display configuration parameters.
6. Under Profile Details, do the following:
   a. Make sure **Virtual AP** enable is selected.
   b. From the **VLAN** drop-down menu, select the VLAN ID to use for the Virtual AP profile.
   c. From the **Forward mode** drop-down menu, select **bridge**.
   d. From the **Remote-AP Operation** drop-down menu, select **always** or **backup**.
   e. Click **Apply**.
7. Under **Profiles**, select **AP**, then **AP system profile**.
8. Under **Profile Details**, do the following:
   a. Select the AP system profile to edit.
   b. At the **LMS IP** field, enter the LMS IP address.
   c. At the **Master controller IP address** field, enter the master controller IP address.
   d. Configure the **Remote-AP DHCP Server** fields.
   e. Click **Apply**.

Configuring the Session ACL in the CLI

Use the following commands:

```
ip access-list session <policy>
   any any svc-dhcp permit
   any any any route src-nat
```

If you use a local DHCP server to obtain IP addresses, you must define one additional ACL to permit traffic between clients without source NATing the traffic. Add **user alias internal-network any permit** before **any any any route src-nat:**

```
user-role <role>
   session-acl <policy>
```

Using the CLI to configure the AAA profile

Use the following commands:

```
aaa profile <name>
   initial-role <role>
```

You can define other parameters as needed.

Defining the Backup Configuration in the CLI

Use the following commands:

```
wlan ssid-profile <profile>
   essid <name>
   opmode <method>
   wpa-passphrase <string> (if necessary)

wlan virtual-ap <name>
   ssid-profile <profile>
   vlan <vlan>
   forward-mode bridge
   aaa-profile <name>
   rap-operation {always|backup}

ap system-profile <name>
   lms-ip <ipaddr>
   master-ip <ipaddr>
   rap-dhcp-default-router <ipaddr>
   rap-dhcp-dns-server <ipaddr>
   rap-dhcp-lease <days>
   rap-dhcp-pool-end <ipaddr>
   rap-dhacp-pool-netmask <netmask>
   rap-dhcp-pool-start <ipaddr>
   rap-dhcp-server-id <ipaddr>
   rap-dhcp-server-vlan <vlan>

ap-group <name>
   virtual-ap <name>
   ap-system-profile <name>
```

or

```
ap-name <name>
   virtual-ap <name>
   ap-system-profile <name>
```

Specifying the DNS Controller Setting

In addition to specifying IP addresses for controllers, you can also specify the master DNS name for the controller when provisioning the remote AP. The name must be resolved to an IP address when attempting to set up the IPSec tunnel. For information on how to configure a host name entry on the DNS server, refer to the vendor documentation for your server. It is recommended to use a maximum of 8 IP addresses to resolve a controller name.

If the remote AP gets multiple IP addresses responding to a host name lookup, the remote AP can use one of them to establish a connection to the controller. For more detailed information, see the next section Backup Controller List on page 653.

Specifying the name also lets you move or change remote AP concentrators without reprovisioning your APs. For example, in a DNS load-balancing model, the host name resolves to a different IP address depending on the location of the user. This allows the remote AP to contact the controller to which it is geographically closest.

The DNS setting is part of provisioning the AP. The easiest way to provision an AP is to use the Provisioning page in the WebUI. These instructions assume you are only modifying the controller information in the Master Discovery section of the Provision page.



Reprovisioning the AP causes it to automatically reboot.

In the WebUI

1. Navigate to the **Configuration** > **Wireless** > **AP Installation** > **Provisioning** page. Select the remote AP and click **Provision**.
2. Under **Master Discovery** enter the master DNS name of the controller.
3. Click **Apply and Reboot**.

For more information, see Provision the AP on page 635.

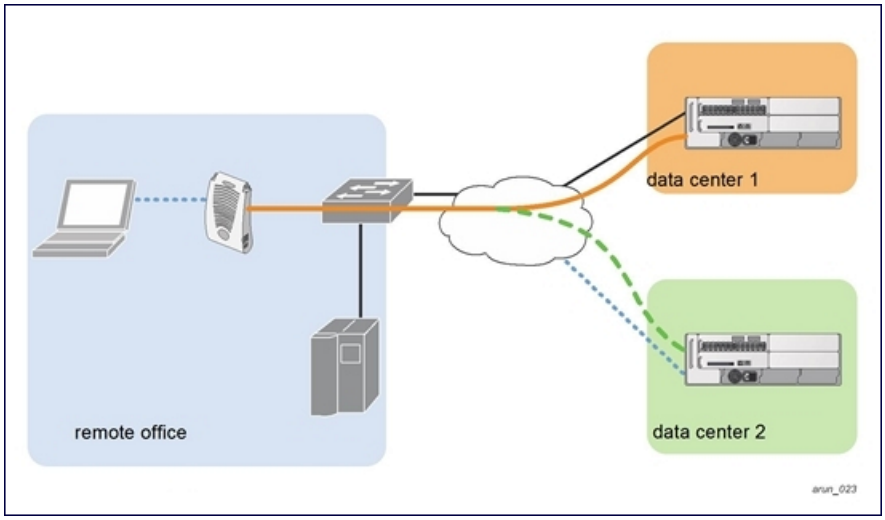
Backup Controller List

Using DNS, the remote AP receives multiple IP addresses in response to a host name lookup. Known as the backup controller list, remote APs go through this list to associate with a controller. If the primary controller is unavailable or does not respond, the remote AP continues through the list until it finds an available controller. This provides redundancy and failover protection.

The remote AP loses the IP address information received through DNS when it terminates and receives the system profile configuration from the controller. If the remote AP loses connectivity on the IPSec tunnel to the controller, the RAP fails over from the primary controller to the backup controller. For this scenario, add the IP address of the backup controller in the backup LMS and the IP address of the primary controller in the LMS field of the ap-system profile. Network connectivity is lost during this time. As described in the section Configuring Remote AP Failback on page 654, you can also configure a remote AP to revert back to the primary controller when it becomes available. To complete this scenario, you must also configure the LMS IP address and the backup LMS IP address.

For example, assume you have two data centers, data center 1 and data center 2, and each data center has one master controller in the DMZ. You can provision the remote APs to use the controller in data center 1 as the primary controller, and the controller in data center 2 as the backup controller. If the remote AP loses connectivity to the primary, it will attempt to establish connectivity to the backup. You define the LMS parameters in the AP system profile.

Figure 78 Sample Backup Controller Scenario



Configuring the LMS and backup LMS IP addresses in the WebUI

1. Navigate to the **Configuration** > **Wireless** > **AP Configuration** page.
2. Select either the **AP Group** or **AP Specific** tab. Click **Edit** for the AP group or AP name.
3. Under Profiles, select **AP** to display the AP profiles.
4. Select the AP system profile you want to modify.

5. Under **Profile Details**:
   a. At the **LMS IP** field, enter the primary controller IP address.
   b. At the **Backup LMS IP** field, enter the backup controller IP address.
6. Click **Apply**.

Configuring the LMS and backup LMS IP addresses in the CLI

Use the following commands:

```
ap system-profile <profile>
   lms-ip <ipaddr>
   bkup-lms-ip <ipaddr>

ap-group <group>
   ap-system-profile <profile>

ap-name <name>
   ap-system-profile <profile>
```

Configuring Remote AP Failback

In conjunction with the backup controller list, you can configure remote APs to revert back (failback) to the primary controller if it becomes available. If you do not explicitly configure this behavior, the remote AP will keep its connection with the backup controller until the remote AP, controller, or both have rebooted or some type of network failure occurs. If any of these events occur, the remote AP will go through the backup controller list and attempt to connect with the primary controller.

In the WebUI

1. Navigate to the **Configuration** > **Wireless** > **AP Configuration** page.
2. Select either the AP Group or AP Specific tab. Click **Edit** for the AP group or AP name.
3. Under **Profiles**, select **AP** to display the AP profiles.
4. Select the AP system profile you want to modify.
5. Under **Profile Details**:
   a. Click **LMS Preemption**. This is disabled by default.
   b. At the **LMS Hold-down period** field, enter the amount of time the remote AP must wait before moving back to the primary controller.
6. Click **Apply**.

In the CLI

Use the following commands:

```
ap system-profile <profile>
   lms-preemption
   lms-hold-down period <seconds>
```

Enabling RAP Local Network Access

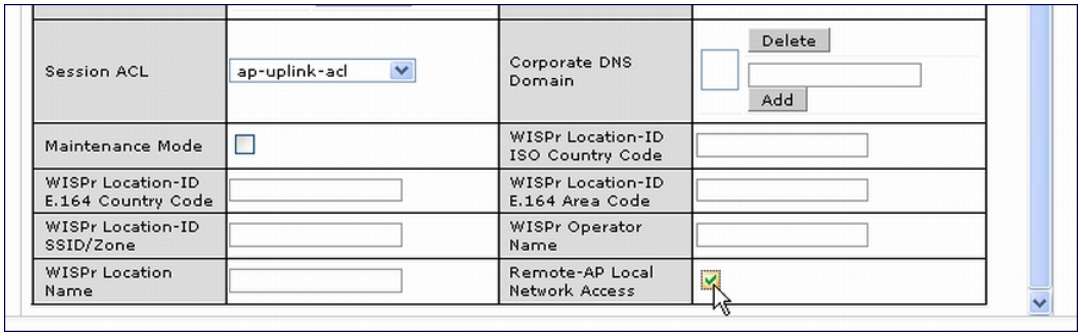
You can enable local network access between the clients (from same or different subnets and VLANs) connected to a RAP through wired or wireless interfaces in split-tunnel/bridge forwarding modes. This allows the clients to effectively communicate with each other without routing the traffic via the controller. You can use WebUI or CLI to enable the local network access.

In the WebUI

1. Navigate to the **Configuration** > **Wireless** > **AP Configuration** page.

2. Select the **AP Group** tab. Click **Edit** for the AP group or AP name.
3. Under **Profiles**, expand the **AP menu,** then select **AP system profile**.
4. To enable remote network access, select the **Remote-AP Local Network Access** check box.

Figure 79 *Enable Remote AP Local Network Access*



5. Click **Apply**.

In the CLI

- To enable, enter the following command:

```
ap system-profile <ap-profile> rap-local-network-access
```

- To disable, enter the following command:

```
ap system-profile <ap-profile> no rap-local-network-access
```

See the *Dell Networking W-Series ArubaOS Command Line Reference Guide* for detailed information on the command options.

Configuring Remote AP Authorization Profiles

Remote AP configurations include an authorization profile that specifies which profile settings should be assigned to a remote AP that has been provisioned but not yet authenticated at the remote site. These yet-unauthorized APs are put into the temporary AP group **authorization-group** by default and assigned the predefined profile **NoAuthApGroup**. This configuration allows the user to connect to an unauthorized remote AP via a wired port, then enter a corporate username and password. Once a valid user has authorized the AP, and it will be marked as authorized on the network. The remote AP will then download the configuration assigned to that AP by its permanent AP group.

In the WebUI

Adding or Editing a Remote AP Authorization Profile

To create a new authorization profile or edit an existing authorization profile via the WebUI:

1. Select **Configuration** > **All Profiles**. The **All Profile Management** window opens.
2. Select **AP** to expand the **AP** profile menu.
3. Select **AP Authorization Profile**. The **Profile Details** pane appears and displays the list of existing AP authorization profiles.
   - To edit an existing profile, select a profile from from the **Profile Details** pane.
   - To create a new authorization profile, enter a new profile name in the entry blank on the **Profile Details** pane, then click **Add**.
4. The **Profile Details** window will display the AP group currently defined for that authorization profile. To select a new AP group, click the drop-down list and select a different AP group name.

5. Click **Apply**.

In the CLI

To create a new authorization profile or edit an existing authorization profile via the command-line interface, access the command-line interface in enable mode, and issue the following commands.

```
ap authorization-profile <profile>
   authorization-group <ap-group>
```

Working with Access Control Lists and Firewall Policies

Remote APs support the following access control lists (ACLs); unless otherwise noted, you apply these ACLS to user roles:

- Standard ACLs—Permit or deny traffic based on the source IP address of the packet.
- Ethertype ACLs—Filter traffic based on the Ethertype field in the frame header.
- MAC ACLs—Filter traffic on a specific source MAC address or range of MAC addresses.
- Firewall policies (session ACLs)—Identifies specific characteristics about a data packet passing through the Dell controller and takes some action based on that identification. You apply these ACLs to user roles or uplink ports.



To configure firewall policies, you must install the PEFNG license.

For more information about ACLs and firewall policies, see Configuring Fallback Mode on page 645.

Understanding Split Tunneling

The split tunneling feature allows you to optimize traffic flow by directing only corporate traffic back to the controller, while local application traffic remains local. This ensures that local traffic does not incur the overhead of the round trip to the controller, which decreases traffic on the WAN link and minimizes latency for local application traffic. This is useful for sites that have local servers and printers. With split tunneling, a remote user associates with a single SSID, not multiple SSIDs, to access corporate resources (for example, a mail server) and local resources (for example, a local printer). The remote AP examines session ACLs to distinguish between corporate traffic destined for the controller and local traffic.

Figure 80 Sample Split Tunnel Environment

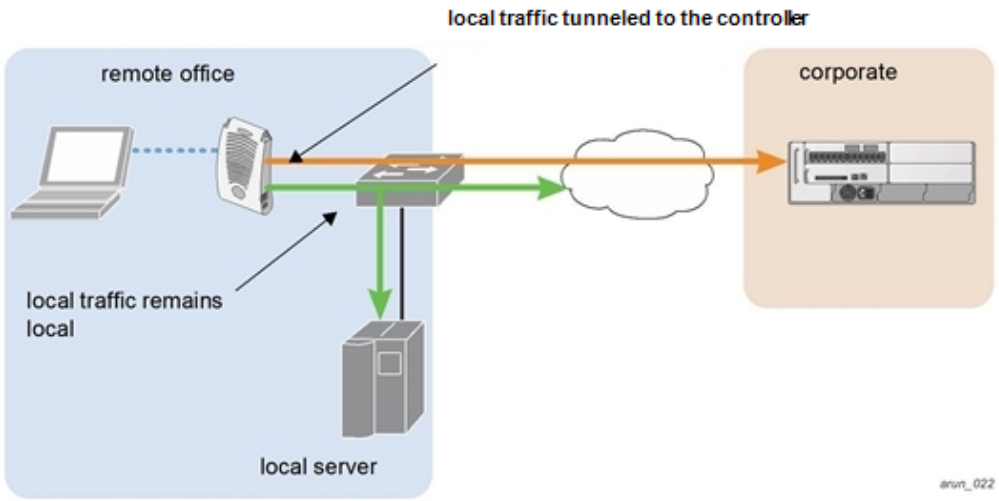


Figure 80 displays corporate traffic is GRE tunneled to the controller through a trusted tunnel and local traffic is source NATed and bridged on the wired interface based on the configured user role and session ACL.

Configuring Split Tunneling

The procedure to configure split tunneling requires the following steps. Each step is described in detail later in this chapter.



The split tunneling feature requires the PEFNG license. If you do not have the PEFNG license on your controller, you must install it before you configure split tunneling. For details on installing licenses, see Software Licenses on page 130.

1. Define a session ACL that forwards only corporate traffic to the controller.
   a. Configure a net destination for the corporate subnets.
   b. Create rules to permit DHCP and corporate traffic to the corporate controller.
   c. Apply the session ACL to a user role. For information about user roles and policies, see Roles and Policies on page 364.
2. (Optional) Configure an ACL that restricts remote AP users from accessing the remote AP local debugging homepage.
3. Configure the remote AP's AAA profile.
   a. Specify the authentication method (**802.1x** or **PSK**) and the default user role for authenticated users. The user role specified in the AAA profile must contain the session ACL defined in the previous step.
   b. (Optional) Use the remote AP's AAA profile to enable RADIUS accounting.
4. Configure the virtual AP profile:
   a. Specify which AP group or AP to which the virtual AP profile applies.
   b. set the VLAN used for split tunneling. Only one VLAN can be configured for split tunneling; VLAN pooling is not allowed.
   c. When specifying the use of a split tunnel configuration, use "split-tunnel" forward mode.
   d. Create and apply the applicable SSID profile.



When creating a new virtual AP profile In the WebUI, you can also configure the SSID at the same time. For information about AP profiles, see Understanding AP Configuration Profiles on page 488.

5. (Optional) Create a list of network names resolved by corporate DNS servers.

Configuring the Session ACL Allowing Tunneling

First you need to configure a session ACL that "permits" corporate traffic to be forwarded (tunneled) to the controller, and that routes, or locally bridges, local traffic.

Using the WebUI

1. Navigate to the **Configuration** > **Security** > **Access Control** > **Policies** page.
2. Click **Add** to create a new policy.
3. Enter the policy name in the **Policy Name** field.
4. From the **Policy Type** drop-down list, select **Session**.
5. From the **IP Version** drop-down list, select **IPv4** or **IPv6**.
6. To create the first rule:
   a. Under Rules, click **Add**.

b. Under Source, select **any**.
c. Under Destination, select **any**.
d. Under Service, select **service**. In the service drop-down list, select **svc-dhcp**.
e. Under Action, select **permit**forIPv4 or**captive**for IPv6.
f. Click **Add**.

7. To create the next rule:
   a. Under Rules, click **Add**.
   b. Under Source, select **any**.
   c. Under Destination, select **alias**.

   The following steps define an alias representing the corporate network. Once defined, you can use the alias for other rules and policies. You can also create multiple destinations the same way.
8. Under the alias section, click **New**. Enter a name in the Destination Name field.
   a. Click **Add**.
   b. For Rule Type, select **Network**.
   c. Enter the public IP address of the controller.
   d. Enter the Network Mask/Range.
   e. Click **Add** to add the network range.
   f. Click **Apply**. The new alias appears in the Destination menu.
9. Under **Destination**, select the alias you just created.
10. Under **Service**, select **any**.
11. Under **Action**, select **permit**for IPv4 or **captive**for IPv6.
12. Click **Add**.
13. To create the next rule:
    a. Under **Rules**, click **Add**.
    b. Under **Source**, select **user**.
    c. Under **Destination**, select **any**.
    d. Under **Service**, select **any**.
    e. Under **Action**, select **any** and check **src-nat**.
    f. Click **Add**.
14. Click **Apply**.
15. Click the **User Roles** tab.
    a. Click **Add** to create and configure a new user role.
    b. Enter the desired name for the role in the **Role Name** field.
    c. Under **Firewall Policies**, click **Add**.
    d. From the **Choose from Configured Policies** drop-down menu, select the policy you just configured.
    e. Click **Done**.
16. Click **Apply**.

Using the CLI

Use the following commands:

```
ap system-profile <profile>
   lms-preemption
   lms-hold-down period <seconds>netdestination <policy>
   network <ipaddr> <netmask>
```

```
   network <ipaddr> <netmask>

ip access-list session <policy>
   any any svc-dhcp permit
   any alias <name> any permit
   user any any route src-nat

user-role <role>
   session-acl <policy>
```

When defining the alias, there are a number of other session ACLs that you can create to define the handling of local traffic, such as:

```
ip access-list session <policy>
   user alias <name> any redirect 0
   user alias <name> any route
   user alias <name> any route src-nat
```

Configuring an ACL to Restrict Local Debug Homepage Access

A user in split or bridge role using a remote AP (RAP) can log on to the local debug (LD) homepage and perform a reboot or reset operations. The LD homepage provides various information about the RAP and also has a button to reboot the RAP. You can now restrict a RAP user from resetting or rebooting a RAP by using the `localip` keyword in the in the user role ACL.



You will require the PEFNG license to use this feature. See Software Licenses on page 130 for more information on licensing requirements.

Any user associated to that role can be allowed or denied access to the LD homepage. You can use the `localip` keyword in the ACL rule to identify the local IP address on the RAP. The `localip` keyword identifies the set of all local IP addresses on the system to which the ACL is applied. The existing keywords controller and `mswitch` indicate only the primary IP address on the controller.

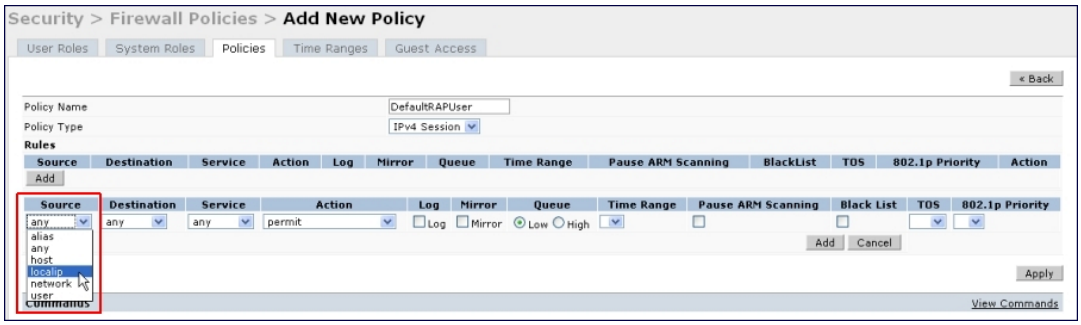


This release of ArubaOS provides localip keyword support only for RAP and not for controller.

In the WebUI

1. Navigate to the **Configuration** > **Security** > **Access Control** > **Policies** page.
2. Click **Add** to create a new policy.
3. Enter the policy name in the **Policy Name** field.
4. From the **Policy Type** drop-down list, select **IPv4 Session**.
5. To create the first rule:
   a. Under **Rules**, click **Add**.
   b. Under **Source**, select **localip**.
   c. Under **Destination**, select **any**.
   d. Under **Action**, select **permit**.
   e. Click **Apply**.

Figure 81 *Enable Restricted Access to LD Homepage*



In the CLI

Use the `localip` keyword in the user role ACL.

All users have an ACL entry of type `any any deny` by default. This rule restricts access to all users. When the ACL is configured for a user role, if a `user any permit` ACL rule is configured, add a deny ACL before that for `localip` for restricting the user from accessing the LD homepage.

Example:

```
ip access-list session logon-control
   user localip svc-http deny
   user any permit
```

Configuring the AAA Profile for Tunneling

After you configure the session ACL, you define the AAA profile used for split tunneling. When defining the AAA parameters, specify the previously configured user role that contains the session ACL used for split tunneling.

If you enable RADIUS accounting in the AAA profile, the controller sends a RADIUS accounting start record to the RADIUS server when a user associates with the remote AP, and sends a stop record when the user logs out or is deleted from the user database. If you enable interim accounting, the controller sends updates at regular intervals. Each interim record includes cumulative user statistics, including received bytes and packets counters. For more information on RADIUS accounting, see RADIUS Accounting on page 245

In the WebUI

1. Navigate to the **Security** > **Authentication** > **AAA Profiles** page. From the AAA Profiles Summary list, click **Add**.
2. Enter the AAA profile name, then click **Add**.
3. Select the AAA profile that you just created.
   a. For **802.1X Authentication Default Role**, select the user role you previously configured for split tunneling, then click **Apply**.
   b. Under the AAA profile that you created, locate **802.1x Authentication Server Group**, and select the authentication server group to use, then click **Apply**.
4. (Optional) To enable RADIUS accounting:
   a. Select the AAA profile from the profile list to display the list of authentication and accounting profiles associated with the AAA profile.
   b. Select the **Radius Accounting Server Group** profile associated with the AAA profile. Click the **RADIUS Accounting Server Group** drop-down list to select a RADIUS server group. (For more information on configuring a RADIUS server or server group, see Configuring a RADIUS Server on page 226.)

c. To enable RADIUS Interim Accounting, select the AAA profile name from the profile list, then click the **RADIUS Interim Accounting** checkbox. This option is disabled by default, allowing the controller to send only *start* and *stop* messages to the RADIUS accounting server.

5. Click **Apply**.

If you need to create an authentication server group, select **new** and enter the appropriate parameters.

Inthe CLI

Use the following commands:

```
aaa profile <name>
   authentication-dot1x <dot1x-profile>
   dot1x-default-role <role>
   dot1x-server-group <group>
   radius-accounting <group>
   radius-interim-accounting
```

Configuring the Virtual AP Profile

In the WebUI

1. Navigate to **Configuration** > **Wireless** > **AP Configuration** page. Select either the **AP Group** or **AP Specific** tab. Click **Edit** for the applicable AP group name or AP name.
2. Under **Profiles**, select **Wireless LAN**, then **Virtual AP**.
3. To create a new virtual AP profile in the WebUI, select **New** from the **Add a profile** drop-down menu. Enter the name for the virtual AP profile, and click **Add**.



Whenever you create a new virtual AP profile in the WebUI, the profile automatically contains the “default” SSID profile with the default ESSID. You must configure a new ESSID and SSID profile for the virtual AP profile before you apply the profile.

   a. In the **Profile Details** entry, go to the AAA Profile drop-down list and select the previously configured AAA profile. The AAA Profile pop-up window appears.
   b. To set the AAA profile and close the window, click **Apply**.
   c. In the **Profile Details** entry for the new virtual AP profile, select **NEW** from the **SSID Profile** drop-down menu. A pop-up window displays to allow you to configure the SSID profile.
   d. Enter the name for the SSID profile.
   e. Under **Network**, enter a name in the Network Name (SSID) field.
   f. Under **Security**, select the network authentication and encryption methods.
   g. To set the SSID profile and close the window, click **Apply**.
4. Click **Apply** at the bottom of the Profile Details window.
5. Click the new virtual AP name in the **Profiles** list or the **Profile Details** to display configuration parameters.
6. Under **Profile Details**:
   a. Make sure **Virtual AP enable** is selected.
   b. From the **VLAN** drop-down menu, select the VLAN ID for the VLAN to be used for split tunneling.
   c. From the **Forward mode** drop-down menu, select **split-tunnel**.
   d. Click **Apply**.

In the CLI

Use the following commands:

```
wlan ssid-profile <profile>
```

```
      essid <name>
      opmode <method>

   wlan virtual-ap <profile>
      ssid-profile <name>
      forward-mode <mode>

      vlan <vlan id>
      aaa-profile <profile>

   ap-group <name>
      virtual-ap <profile>
```

or

```
   ap-name <name>
      virtual-ap <profile>
```

Defining Corporate DNS Servers

Clients send DNS requests to the corporate DNS server address that it learned from DHCP. If configured for split tunneling, corporate domains and traffic destined for corporate use the corporate DNS server. For non-corporate domains and local traffic, other DNS servers can be used.

In the WebUI

1. Navigate to **Configuration > Wireless > AP Configuration** page.
2. Select either the **AP Group** or **AP Specific** tab. Click **Edit** for the AP group or AP name.
3. Under Profiles, select **AP**, then **AP system profile**.
4. Under **Profile Details**:
   a. Enter the corporate DNS servers.
   b. Click **Add**.
      The DNS name appears in Corporate DNS Domain list. You can add multiple names the same way.
5. Click **Apply**.

In the CLI

Use the following commands:

```
ap system-profile <profile>
   dns-domain <domain name>
```

Understanding Bridge

The bridge feature allows you to route the traffic flow only to the internet and not to the corporate network. Only the 802.1X authentication request is sent to the corporate network. This feature is useful for guest users.



ArubaOS does not support Wired 802.1x authentication in bridge mode for RAP and CAP. 802.1x authentication is supported only in tunnel and split modes.

Figure 82 Sample Bridge Environment

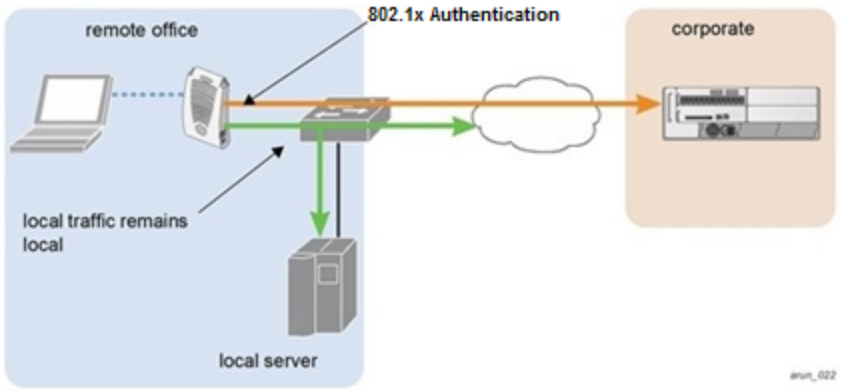


Figure 82 displays the local traffic being routed to the internet and the 802.1X authentication request sent to the corporate network.

Configuring Bridge

To configure a bridge, perform the following steps. Each step is described in detail later in this chapter.

NOTE

The bridge feature requires the PEFNG license. If you do not have the PEFNG license on your controller, you must install it before you configure bridge. For details on installing licenses, see Software Licenses on page 130.

1. Define a session ACL that routes the traffic.
   a. Create rules to permit DHCP and local data traffic.
   b. Apply the session ACL to a user role. For information about user roles and policies, see Roles and Policies on page 364.
2. Configure the remote AP's AAA profile.
   a. Specify the authentication method (**802.1x** or **PSK**) and the default user role for authenticated users. The user role specified in the AAA profile must contain the session ACL defined in the previous step.
   b. (Optional) Use the remote AP's AAA profile to enable RADIUS accounting.
3. Configure the virtual AP profile:
   a. Specify the AP group or ap-name to which the virtual AP profile applies.
   b. Set the VLAN in the virtual AP.
   c. When specifying the use of a bridge configuration, use bridge forward mode.
   d. Create and apply the applicable SSID profile.
   e. (Optional) Under AP system profile, configure the RAP DHCP pool. RAP DHCP VLAN must be same as VAP's VLAN. If the client needs to obtain from the RAP DHCP Server.

NOTE

When creating a new virtual AP profile In the WebUI, you can simultaneously configure the SSID. For information about AP profiles, see Understanding AP Configuration Profiles on page 488.

Configuring the Session ACL

First you need to configure a session ACL that "permits" corporate traffic to be forwarded to the controller and that routes, or locally bridges, local traffic.

Using the WebUI

1. Navigate to the **Configuration** > **Security** > **Access Control** > **Policies** page.

2. Click **Add** to create a new policy.
3. Enter the policy name in the **Policy Name** field.
4. From the **Policy Type** drop-down list, select **Session**.
5. From the **IP Version** drop-down list, select **IPv4** or **IPv6**.
6. To create the first rule:
   a. Under **Rules**, click **Add**.
   b. Under **Source**, select **any**.
   c. Under **Destination**, select **any**.
   d. Under **Service**, select **service**. In the service drop-down list, select **svc-dhcp**.
   e. Under **Action**, select **permit** for IPv4 or **captive** for IPv6.
   f. Click **Add**.
7. 7. To create the next rule:
   a. a. Under Rules, click **Add**.
   b. b. Under Source, select **any**.
   c. c. Under Destination, select **alias**.

The following steps define an alias representing the corporate network. Once defined, you can use the alias for other rules and policies. You can also create multiple destinations the same way.

8. Under the alias section, click **New**. Enter a name in the Destination Name field.
   a. Click **Add**.
   b. For Rule Type, select **Network**.
   c. Enter the public IP address of the controller.
   d. Enter the Network Mask/Range.
   e. Click **Add** to add the network range.
   f. Click **Apply**. The new alias appears in the Destination menu.
9. Under **Destination**, select the alias you just created.
10. Under **Service**, select **any**.
11. Under **Action**, select **permit** for IPv4 or **captive** for IPv6.
12. Click **Add**.
13. To create the next rule:
   a. Under **Rules**, click **Add**.
   b. Under **Source**, select **user**.
   c. Under **Destination**, select **any**.
   d. Under **Service**, select **any**.
   e. Under *Action*, select **any** and check **src-nat**.
   f. Click **Add**.
14. Click **Apply**.
15. Click the **User Roles** tab.
   a. Click **Add** to create and configure a new user role.
   b. Enter the desired name for the role in the **Role Name** field.
   c. Under **Firewall Policies**, click **Add**.
   d. From the **Choose from Configured Policies** drop-down menu, select the policy you just configured.
   e. Click **Done**.

16. Click **Apply**.

Using the CLI

If `dhcp server` in `ap system profile` is enabled

```
ip access-list session <policy> any any svc-dhcp permit
user any any route src-nat
```

If `dhcp server` in `ap system profile` is disabled

```
ip access-list session <policy>
any any any permit
user-role <role>
   session-acl <policy>
```



To configure an ACL to Restrict Local Debug Homepage Access, see Configuring an ACL to Restrict Local Debug Homepage Access on page 659.

Configuring the AAA Profile for Bridge

After you configure the session ACL, you define the AAA profile used for bridge. When defining the AAA parameters, specify the previously configured user role that contains the session ACL used for bridge.

If you enable RADIUS accounting in the AAA profile, the controller sends a RADIUS accounting start record to the RADIUS server when a user associates with the remote AP, and sends a stop record when the user logs out or is deleted from the user database. If you enable interim accounting, the controller sends updates at regular intervals. Each interim record includes cumulative user statistics, including received bytes and packets counters. For more information on RADIUS accounting, see RADIUS Accounting on page 245.

In the WebUI

1. Navigate to the **Security** > **Authentication** > **AAA Profiles** page. From the AAA Profiles Summary list, click **Add**.
2. Enter the AAA profile name, then click **Add**.
3. Select the AAA profile that you just created.
   a. For **802.1X Authentication Default Role**, select the user role you previously configured for split tunneling or bridge, then click **Apply**.
   b. Under the AAA profile that you created, locate **802.1x Authentication Server Group**, and select the authentication server group to use, then click **Apply**.
4. (Optional) To enable RADIUS accounting:
   a. Select the AAA profile from the profile list to display the list of authentication and accounting profiles associated with the AAA profile.
   b. Select the **Radius Accounting Server Group** profile associated with the AAA profile. Click the **RADIUS Accounting Server Group** drop-down list to select a **RADIUS server group**. (For more information on configuring a RADIUS server or server group, see Configuring a RADIUS Server on page 226.)
   c. To enable **RADIUS Interim Accounting**, select the **AAA profile name** from the profile list, then click the **RADIUS Interim Accounting** checkbox. This option is disabled by default, allowing the controller to send only start and stop messages RADIUS accounting server.
5. Click **Apply**.

If you need to create an authentication server group, select **new** and enter the appropriate parameters.

In the CLI

Use the following command:

```
aaa profile <name>
```

```
authentication-dot1x <dot1x-profile>
dot1x-default-role <role>
dot1x-server-group <group>
radius-accounting <group>
radius-interim-accounting
```

Configuring Virtual AP Profile

In the WebUI

1. Navigate to **Configuration** > **Wireless** > **AP Configuration page**. Select either the **AP Group** or **AP Specific** tab. Click **Edit** for the applicable AP group name or AP name.
2. Under **Profiles**, select **Wireless LAN**, then **Virtual AP**.
3. To create a new virtual AP profile in the WebUI, select **New** from the **Add a profile** drop-down menu. Enter the name for the virtual AP profile, and click **Add**.



Whenever you create a new virtual AP profile in the WebUI, the profile automatically contains the “default” SSID profile with the default ESSID. You must configure a new ESSID and SSID profile for the virtual AP profile before you apply the profile.

   a. In the **Profile Details** entry, go to the **AAA Profile** drop-down list and select the previously configured AAA profile. The AAA Profile pop-up window appears.
   b. To set the AAA profile and close the window, click **Apply**.
   c. In the **Profile Details** entry for the new virtual AP profile, select **NEW** from the **SSID Profile** drop-down menu. A pop-up window displays to allow you to configure the SSID profile.
   d. Enter the name for the SSID profile.
   e. Under **Network**, enter a name in the Network Name (SSID) field.
   f. Under **Security**, select the network authentication and encryption methods.
   g. To set the SSID profile and close the window, click **Apply**.
4. Click **Apply** at the bottom of the Profile Details window.
5. Click the new virtual AP name in the Profiles list or the Profile Details to display configuration parameters.
6. Under **Profile Details**:
   a. Make sure **Virtual AP enable** is selected.
   b. From the **VLAN** drop-down menu, select the VLAN ID for the VLAN to be used for bridge.
   c. From the **Forward mode** drop-down menu, select **Bridge**.
   d. Click **Apply**.

In the CLI

Use the following command:

```
wlan ssid-profile <profile> essid <name>
opmode <method>

wlan virtual-ap <profile>
ssid-profile <name>
forward-mode  bridge
vlan <vlan id>
aaa-profile <profile>

ap-group <name>
virtual-ap <profile>
```

or

```
ap-name <name>
virtual-ap <profile>
```

Provisioning Wi-Fi Multimedia

Wi-Fi Multimedia (WMM) is a Wi-Fi Alliance specification based on the IEEE 802.11e wireless Quality of Service (QoS) standard. WMM works with 802.11a, b, g, and n physical layer standards. The IEEE 802.11e standard also defines the mapping between WMM access categories (ACs) and Differentiated Services Codepoint (DSCP) tags. Remote APs support WMM.

WMM supports four ACs: voice, video, best effort, and background. You apply and configure WMM in the SSID profile.

When planning your configuration, make sure that immediate switches or routers do not have conflicting 802.1p or DSCP configurations/mappings. If this occurs, your traffic may not be prioritized correctly.

For more detailed information about WMM and the applicable configuration commands, see Voice and Video on page 878.

Reserving Uplink Bandwidth

You can reserve and prioritize uplink bandwidth traffic to provide higher QoS for specific applications, traffic, or ports. This is done by applying bandwidth reservation on existing session ACLs. Typically, the bandwidth reservation is applied for uplink voice traffic.

Note the following before you configure bandwidth reservation:

- You must know the total bandwidth available.
- The bandwidth reservation are applicable only on session ACLs.
- Bandwidth reservation on voice traffic ACLs receives higher priority over other reserved traffic.
- You can configure up to three unique priority for bandwidth reservation.
- The bandwidth reservation must be specified in absolute value (kbps).
- Priorities for bandwidth reservation are optional, and bandwidth reservations without priorities are treated equal.

Understanding Bandwidth Reservation for Uplink Voice Traffic

The voice ACLs are applicable on the voice signalling traffic used to establish voice call through a firewall. When a voice ACL is executed, a dynamic session is introduced to allow voice traffic through the firewall. This prevents the re-use of voice ACLs for bandwidth reservation. However, you can create bandwidth reservation rules that can be applied on voice signalling traffic and ports used for voice data traffic. This mechanism filters traffic as per the security requirements.

Configuring Bandwidth Reservation

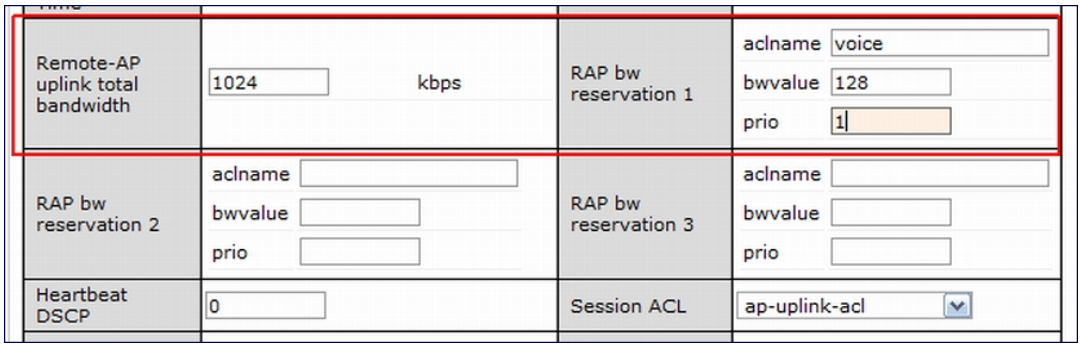
You can configure bandwidth reservation ACLs using the WebUI or the CLI.

In the WebUI

To configure bandwidth reservation

1. Navigate to **Configuration** > **Advanced Services** > **All Profiles**
2. Under **Profiles**, navigate to **AP** > **AP System Profile**. You can create a new AP system profile to configure bandwidth reservation or edit an existing AP system profile. Under the **Profiles Details** page, specify bandwidth reservation values.

Figure 83 *Uplink Bandwidth Reservation*



In the CLI

Use the following commands:

```
(host) (config)#ap system-profile remotebw
(host) (AP system profile "remotebw") #rap-bw-total 1024
(host) (AP system profile "remotebw") #rap-bw-resv-1 acl voice 128 priority 1
```

To view bandwidth reservations:

```
(host) #show datapath rap-bw-resv ap-name remote-ap-1

RAP Uplink BW reservation statistics
------------------------------------
Pos: Acl     Resv Prio XmitPkts XmitByte   Marked Enqueued  Onqueue    Drops TokenFin
-------------------------------------------------------------------------------------
1  :  11      200    0        0        0        3        0        0        0        0
2  :   0        0    0        0        0        0        0        0        0        0
3  :   0        0    0        0        0        0        0        0        0        0
4  :   0        0    0     1524   370962        0     1524        0        0        0
```

Provisioning 4G USB Modems on Remote Access Points

ArubaOS provides support for 4G networks by allowing you to provision 4G USB modems on the RAP. You can also provision the RAP to support both 4G and 3G USB modems. This enables the RAP to choose the available network automatically. 4G takes precedence over 3G when the RAP tries to auto select the network. You can also configure the RAP to work exclusively on a 3G or 4G network. It is recommended that you provision the USB modems for the RAP based on your network requirements.

4G USB Modem Provisioning Best Practices and Exceptions

- RAP does not support dynamic plug-and-play for the 4G USB modems. You must provision a RAP with the 4G USB parameters on the controller manually based on its type and family (4G-WiMAX/4G-LTE).
- When a RAP connects to a 4G network, it appears as a Remote AP (R) and Cellular (C) on the controller.
- For a 3G/4G network switch, using the UML290 modem with the firmware version L0290VWB522F.242 or later is recommended. Using a lower version of the firmware auto-selects the network mode based on the network availability. The latest version allows the RAP to lock the modem in a particular network mode (for example, 3G only).



The 4G-WiMAX family of modems do not support the 3G-4G network switch-over functionality.

ArubaOS 6.3 includes a new method of provisioning multimode USB modems (such as a Verizon UML290, Verizon MC551L, or AT&T 313u) for a remote AP. These changes simplify modem provisioning for both 3G and 4G networks. The modem configuration procedure in ArubaOS 6.2.0.x and earlier versions required that you define a driver for a 3G modem in the USB modem field under the AP provisioning profile, or define a driver for a 4G modem in the 4G USB type field. Starting with ArubaOS 6.3, you can configure drivers for both a 3G or a 4G modem using the USB field, and the 4G USB Type field is deprecated.

Provisioning RAP for USB Modems

To enable 3G/4G network support, you must provision the RAP with the USB parameters on the controller. You can use the WebUI or CLI to provision the USB parameters.

In the WebUI

1. Navigate to the **Configuration > Wireless > AP Installation** page.
2. Select the **Provisioning** tab.
3. Select an AP and click **Provision**.
4. Select the **Yes** option by **Remote AP**.
5. Under **USB Settings**, select the **USB Parameters** check box.
6. Click the **Device** drop-down list and select a USB modem device.
7. Click the **Cellular NW Preferences** drop-down list and select one of the following provisioning options.

Table 124: *Cellular Network Preference Parameters*

Parameter	Description
auto (default)	In this mode, the modem firmware will control the cellular network service selection; so the cellular network service failover and fallback is not interrupted by the remote AP (RAP).
3g_only	Locks the modem to operate only in 3G .
4g_only	Locks the modem to operate only in 4G .
advanced	The RAP controls the cellular network service selection based on an Received Signal Strength Indication (RSSI) threshold-based approach. • Initially the modem is set to the default auto mode. This allows the modem firmware to select the available network. • The RAP determines the RSSI value for the available network type (for example 4G), checks whether the RSSI is within required range, and if so, connects to that network. • If the RSSI for the modem's selected network is not within the required range, the RAP will then check the RSSI limit of an alternate network (for example, 3G), and reconnect to that alternate network. The RAP will repeat the above steps each time it tries to connect using a 4G multimode modem in this mode.

8. Click **Apply and Reboot** to reboot the RAP with the new configuration.

In the CLI

To enable 4G-exclusive network support on the RAP, execute the following commands:

```
(host) (config) #ap provisioning-profile <profile-name>
(host) (Provisioning profile "<profile-name>") usb-type <USB modem type>
(host) (Provisioning profile "<profile-name>") #usb-type none
(host) (Provisioning profile "<profile-name>") #cellular_nw_preference 4g_only
```

To enable 3G-exclusive network support on the RAP, execute the following commands:

```
(host) (config) #ap provisioning-profile <profile-name>
(host) (Provisioning profile "<profile-name>") usb-type <USB modem type>
(host) (Provisioning profile "<profile-name>") #usb-type none
(host) (Provisioning profile "<profile-name>") #cellular_nw_preference 3g_only
```

To enable 3G/4G network switch support, execute the following commands:

```
(host) (config) #ap provisioning-profile <profile-name>
(host) (Provisioning profile "<profile-name>") usb-type <USB modem type>
(host) (Provisioning profile "<profile-name>") #usb-type none
(host) (Provisioning profile "<profile-name>") #cellular_nw_preference auto
```

RAP 3G/4G Backhaul Link Quality Monitoring

The RAP is enhanced to support link monitoring on 2G, 3G, and 4G modems to provide information about the state of USB modem and cellular network.

The USB modem has the following four states:

- **Active** - The USB modem is used as the primary path for connecting VPN to the controller
- **Standby** or **Backup** - The network is available but the USB modem is not used for connecting VPN to the controller
- **Error** - The USB modem is available but the modem is faulty
- **Not Plugged** - The USB modem is unavailable

To view the USB modem details on the RAP, execute the following command:

```
(host) #show ap debug usb ap-name <ap-name>
```

The following is a sample output that shows the USB information provisioned on the RAP:

```
USB Information
---------------
Parameter                   Value
---------                   -----
Manufacturer                 Pantech,
Product                      PANTECH
Serial Number
Driver                       ptuml_cdc_ether
Vendor ID                    106c
Product ID                   3718
USB Modem State              Active
USB Uplink RSSI(in dBm)      -73
Supported Network Services   CDMA GSM LTE
Firmware Version             L0290VWB522F.242
ESN Number                   990000472325325
Current Network Service      4G-LTE
```

Provisioning RAPs at Home

The following section provides information on provisioning your remote AP (RAP) at home using a static IP address, PPPoE connection, or USB modem.

Prerequisites

Follow the steps below to acquire a static IP address before provisioning the RAP at home:

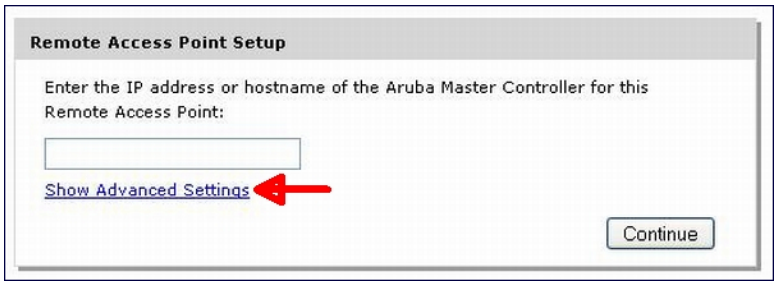
1. Connect the RAP at the site of deployment and ensure that it has connectivity to the Internet to reach the controller.
2. Connect a laptop to Port 1 of the RAP to get an IP address from the RAP's internal DHCP pool.

Provisioning RAP Using Zero Touch Provisioning

You provision the RAP using provisioning wizard:

1. Navigate to the RAP configuration URL : http:/rapconsole.dell-pcw.com.
2. Enter the IP address or hostname of the controller.
3. Click the **Show Advanced Settings** link, shown in Figure 84.

Figure 84 *Show Advanced Settings*



4. In the **Advanced Settings** wizard, you can select one of the following:
   a. **Static IP**—Select this tab to provision your RAP using a static IP address.
   b. **PPPoE**—Select this tab to provision your RAP on a PPPoE connection.
   c. **USB**—Select this tab to provision your RAP using 3G/EVDO USB modem.

Provisioning the RAP using a Static IP Address

Select the **Static IP** tab and enter the required details. See Table 125 for information on parameters.

Figure 85 *Provision RAP using Static IP*

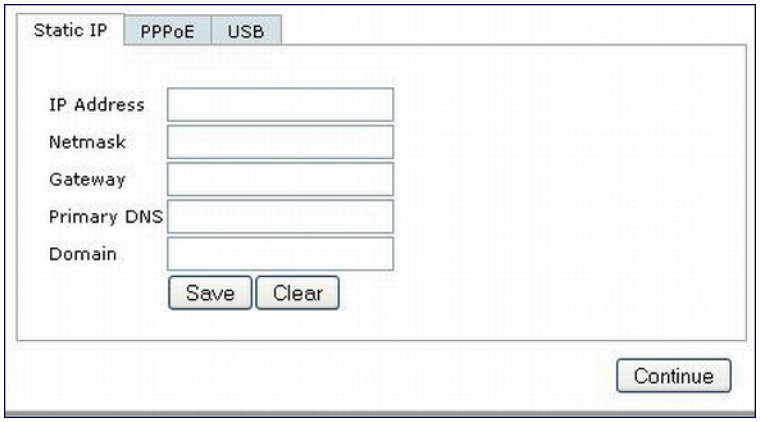


Table 125: *Provision using Static IP*

Item	Description
IP Address	Enter the static IP address that you want to configure for your remote access point.
Netmask	Enter the network mask.
Gateway	Enter the default gateway IP address of your network.

Item	Description
Primary DNS	Enter the IP address of your primary DNS server. This is an optional parameter.
Domain	Enter your domain name. This is an optional parameter.

Click **Save** after you have entered all the details.

Provision the RAP on a PPPoE Connection

Select the **PPPoE** tab and enter the required details. See Table 126 for information on parameters.

Figure 86 *Provision RAP on a PPPoE Connection*

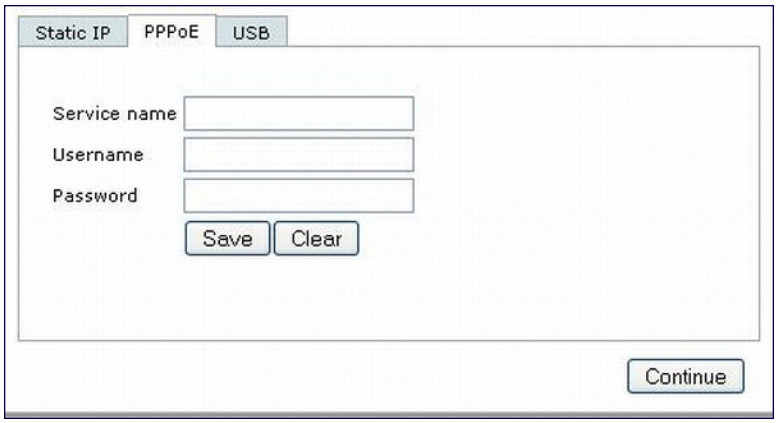


Table 126: *Provision using PPPoE Connection*

Item	Description
Service Name	Enter the PPPoE service name provided to you by your service provider. This parameter is optional.
Username	Enter the user name for the PPPoE connection.
Password	Enter your PPPoE password.

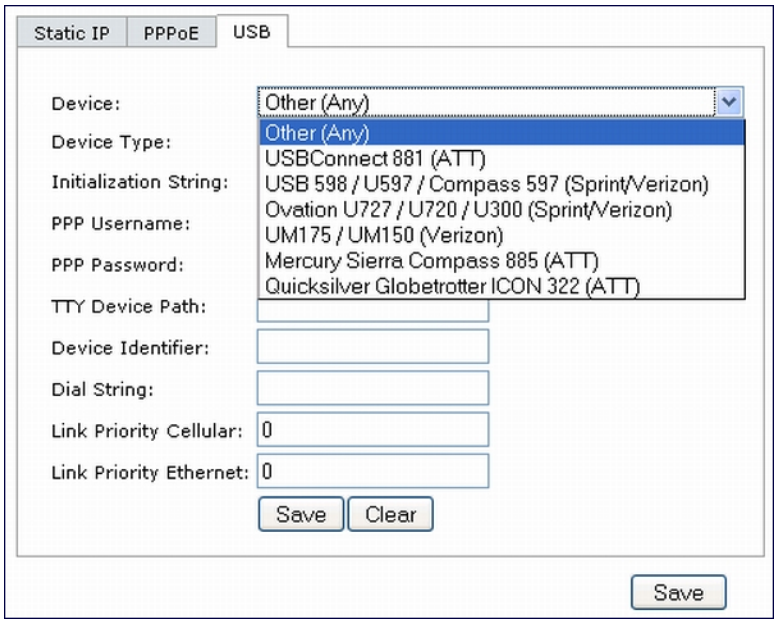
Click **Save** after you have entered all the details.

Using 3G/EVDO USB Modems

The following procedure illustrates provisioning your RAP using a 3G/EVDO USB modem.

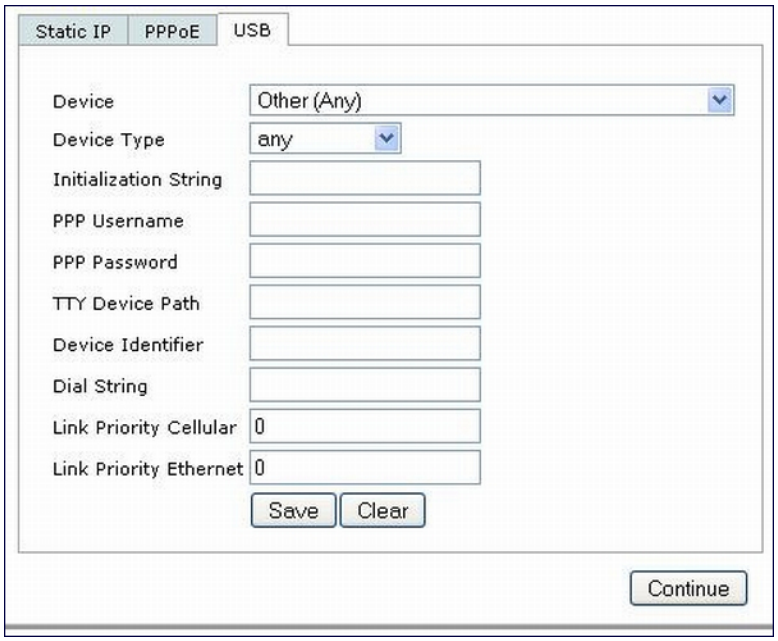
1. Select the **USB** tab and select your modem from the drop down list. Configuration details automatically appear for some common modems.

Figure 87 *Provision using a preconfigured USB Modem*



2. If your modem name is not listed, select **Other** and manually enter the following details. These are available from the manufacturer of your modem or from your IT administrator:

Figure 88 *Provision using a USB Modem with Custom Settings*



- Device Type
- Initializing String
- PPP Username
- PPP Password
- TTY Device Path
- Device Identifier
- Dial String

- Link Priority Cellular—This is a number that identifies the priority of the connection. If the *Link Priority Cellular* has a higher number than *Link Priority Ethernet*, then cellular connection is used.
- Link Priority Ethernet—This is a number that identifies the priority of the connection. If the *Link Priority Ethernet* has a higher number than *Link Priority Cellular*, then Ethernet connection is used.

3. Click **Save** after you have entered all the details and click **Continue** to complete provisioning of your RAP.

Configuring W-IAP3WN and W-IAP3WNP Access Points

The DellW-IAP3WN and W-IAP3WNP are single-radio, single-band wireless APs that support the IEEE 802.11n standard for high-performance WLAN. These APs use MIMO (Multiple-In, Multiple-Out) technology and other high-throughput mode techniques to deliver high-performance, 802.11n 2.4 GHz functionality while simultaneously supporting existing 802.11 b/g wireless services.

See the *Dell Networking W-Series W-IAP3WN/P Installation Guide* for more information.



These access points require Dell Instant 3.0 or later to operate as an Instant AP, or ArubaOS 6.1.4.0 or later to operate as a Remote AP.

The Power Sourcing Equipment (PSE) functionality is available only for W-IAP3WNP APs, as the PoE itself provides the PSE functionality for W-IAP3WN APs. You can use the WebUI or CLI to enable or disable the PSE functionality on the W-IAP3WNP APs.

Using the WebUI

1. Navigate to the **Configuration > Advanced Services > All Profiles** page.
2. Select the **AP** tab, then the **AP Ethernet Link profile** tab.
3. Select the **default** tab .
4. Select the **Power over Ethernet** checkbox.
5. Click **Apply**. Support for W-IAP3WN and W-IAP3WNP access points (APs)

Using the CLI

- To enable, enter:

```
(host)(config) #ap enet-link-profile <name>
  poe
```

- To disable, enter:

```
(host)(config) #ap enet-link-profile <name>
  no poe
```

Use the following command to view the PoE port status on an AP:

```
(host) #show ap enet-link-profile default
  AP Ethernet Link profile "default"
  ----------------------------------
  Parameter Value
  --------- -----
  Speed auto
  Duplex auto
  Power over Ethernet Disabled
```

Converting an IAP to RAP or CAP

For IAP to RAP or CAP conversion, the virtual controller sends the convert command to all the other IAPs. The virtual controller along with the other slave IAPs then set up a VPN tunnel to the remote controller, and

download the firmware by FTP. The Virtual Controller uses IPsec to communicate to the controller over the internet.



A mesh point cannot be converted to RAP because mesh does not support VPN connection.

An IAP can be converted to a Campus AP and Remote AP only if the controller is running ArubaOS 6.1.4 or later.

The following table describes the supported IAP platforms and minimal AOS version for IAP to CAP/RAP conversion.

Converting IAP to RAP

To convert an IAP to RAP, follow the instructions below:

1. Navigate to the **Maintenance** tab in the top right corner of the Instant UI.
2. Click the **Conver**t tab.
3. Select **Remote APs managed by a Mobility Controller** from the drop-down list.
4. Enter the hostname (fully qualified domain name) or the IP address of the controller in the **Hostname or IP Address of Mobility Controller** text box. This information is provided by your network administrator.



Ensure the Mobility Controller IP Address is reachable by the IAPs.

5. Click **Convert Now** to complete the conversion.
6. The IAP reboots and begins operating in RAP mode.
7. After conversion, the IAP is managed by the Dell controller which has been specified in the Instant UI.



In order for the RAP conversion to work, ensure that you configure the Instant AP in the RAP white-list and enable the FTP service on the controller.



If the VPN setup fails and an error message pops up, please click OK, copy the error logs and share them with your Dell support engineer.

Converting an IAP to CAP

To convert an IAP to a Campus AP, do the following:

1. Navigate to the **Maintenance** tab in the top right corner of the Instant UI.
2. Click the **Convert** tab.
3. Select **Campus APs managed by a Mobility Controller** from the drop-down list.
4. Enter the hostname (fully qualified domain name) or the IP address of the controller in the **Hostname or IP Address of Mobility Controller** text box. This is provided by your network administrator.



Ensure the Mobility Controller IP Address is reachable by the IAPs.

5. Click **Convert Now** to complete the conversion.

Enabling Bandwidth Contract Support for RAPs

This release of ArubaOS provides Bandwidth Contract support on remote APs. This is achieved by extending the Bandwidth Contract support on split-tunnel and bridge modes.

You can apply Bandwidth Contract for a RAP on a per-user or per-role basis. Bandwidth Contract is applied on a per-role basis by default. This implies that all the users belonging to the same role will share the bandwidth pool. When Bandwidth Contract configured on the controller is attached to a user-role, it automatically gets pushed to the RAPs terminating on it.

The following show commands have been enhanced in this release to retrieve the Bandwidth Contract information from the RAP:

```
show datapath user ap-name <ap-name>
show datapath bwm ap-name <ap-name>
```

Configuring Bandwidth Contracts for RAP

You can configure bandwidth contracts for RAP on a per-role or per-user basis. The following examples illustrate how to configure, apply, and verify the Bandwidth Contracts on the RAPs.

Defining Bandwidth Contracts

Use the following command to define a 256 Kbps contract:

```
(host) (config) #aaa bandwidth-contract 256k kbits 256
```

Use the following command to define a 512 Kbps contract

```
(host) (config) #aaa bandwidth-contract 512k kbits 512
```

Applying Contracts

You can apply the contract on a per-role or per-user basis.

Applying Contracts Per-Role

Use the following commands to apply the contracts on a per-role basis for upstream and downstream:

For upstream contract of 512 Kbps:

```
(host) (config) #user-role authenticated bw-contract 512k upstream
```

For downstream contract of 256 Kbps:

```
(host) (config) #user-role authenticated bw-contract 256k downstream
```

Applying Contracts Per-User

Use the following commands to apply the contracts on a per-user basis for upstream and downstream:

For upstream contract of 512 Kbps:

```
(host) (config) #user-role authenticated bw-contract 512k per-user upstream
```

For downstream contract of 256 Kbps:

```
(host) (config) #user-role authenticated bw-contract 256k per-user downstream
```

Verifying Contracts on AP

The following example displays the bandwidth contracts on AP for per-role configuration:

```
(host) #show datapath bwm ap-name rap5-2

Datapath Bandwidth Management Table Entries
-------------------------------------------
Contract Types :
0 - CP Dos 1 - Configured contracts 2 - Internal contracts
```

```
-------------------------------------------
Flags: Q - No drop, P - No shape(Only Policed),
T - Auto tuned
---  --------  ---------  ----------  -----------  -----------------
Cont                       Avail   Queued/Pkts
Type  Id   Bits/sec  Policed    Bytes    Bytes      Flags
---- ---- --------- ---------- ------- ------------ -----
1    1      512000          0   16000       0/0    P
1    2      256000          0    8000       0/0    P
```

The following example displays the bandwidth contracts on AP for per-user configuration (contract IDs 3 and 4 are per-user contracts):

```
(host) #show datapath bwm ap-name rap5-2

Datapath Bandwidth Management Table Entries
-------------------------------------------
Contract Types :
0 - CP Dos 1 - Configured contracts 2 - Internal contracts
-------------------------------------------
Flags: Q - No drop, P - No shape(Only Policed),
T - Auto tuned
---  --------  ---------  ----------  -----------  -----------------
Cont                       Avail   Queued/Pkts
Type  Id   Bits/sec  Policed    Bytes    Bytes      Flags
---- ---- --------- ---------- ------- ------------ -----
1    1       512000        300   16000       0/0    P
1    2       256000        277    8000       0/0    P
1    3       512000          0   16000       0/0    P
1    4       256000          0    8000       0/0    P
```

Verifying Contracts Applied to Users

You can verify if the contracts are applied to the user after the user connects to the AP using CLI.

The following is a sample output for a per-role configuration:

```
(host) #show datapath user ap-name rap5-2

Datapath User Table Entries
---------------------------
Flags: P - Permanent, W - WEP, T- TKIP, A - AESCCM, G - AESGCM, V - ProxyArp to/for MN
(Visitor),
N - VPN, L - local, Y - Any IP user, R - Routed user, M - Media Capable,
S - Src NAT with VLAN IP, E - L2 Enforced, F - IPIP Force Delete, O - VOIP user
FM(Forward Mode): S - Split, B - Bridge, N - N/A
IP              MAC           ACLs    Contract   Location  Age  Sessions   Flags   Vlan   FM
--------------  ----------------  -------  ---------  --------  ---  ---------  -----   ----
  --
10.15.72.50     00:0B:86:61:12:AC  2703/0      0/0    0         16       1/65535     P      0
  N
10.15.72.253    00:18:8B:A9:A8:DF    52/0      1/2    0         1        0/65535            1
  S
192.168.11.1    00:0B:86:66:03:3F  2700/0      0/0    0         20024    0/65535     P    177
  N
10.15.196.249   00:0B:86:66:03:3F  2700/0      0/0    0         3        1/65535     P      1
  N
```

The following is a sample output for a per-user configuration:

```
(host) #show datapath user ap-name rap5-2

Datapath User Table Entries
---------------------------
```

```
Flags: P - Permanent, W - WEP, T- TKIP, A - AESCCM, G - AESGCM, V - ProxyArp to/for MN
(Visitor),
N - VPN, L - local, Y - Any IP user, R - Routed user, M - Media Capable,
S - Src NAT with VLAN IP, E - L2 Enforced, F - IPIP Force Delete, O - VOIP user
FM(Forward Mode): S - Split, B - Bridge, N - N/A
IP              MAC           ACLs    Contract   Location  Age  Sessions   Flags   Vlan    FM
--------------  ----------------  -------  ---------  --------  ---  ---------  -----   ----
  --
10.15.72.50     00:0B:86:61:12:AC  2703/0      0/0    0         11       0/65535  P         0
   N
10.15.72.253    00:18:8B:A9:A8:DF    52/0      3/4    0         46       0/65535             1
   S
192.168.11.1    00:0B:86:66:03:3F  2700/0      0/0    0         20883    0/65535  P       177
   N
10.15.196.249   00:0B:86:66:03:3F  2700/0      0/0    0         15       1/65535  P         1
   N
```

Verifying Bandwidth Contracts During Data Transfer

You can verify the Bandwidth Contracts that are in use during data transfer using CLI.

The following is a sample output for a per-role configuration:

```
(host) #show datapath session ap-name rap5-2 table 10.15.72.99

Datapath Session Table Entries
------------------------------
Flags: F - fast age, S - src NAT, N - dest NAT
D - deny, R - redirect, Y - no syn
H - high prio, P - set prio, T - set ToS
C - client, M - mirror, V - VOIP
Q - Real-Time Quality analysis
I - Deep inspect, U - Locally destined
E - Media Deep Inspect, G - media signal
RAP Flags: 1 - Class 1, 2 - Class 2, 3 - Class 3
Source IP     Destination IP  Prot SPort DPort  Cntr Prio ToS Age Destination TAge Flags
--------------  --------------  ---- ----- -----  ---- ---- --- --- ----------- ---- -----
10.15.72.253    10.15.72.99     6    5001  36092  1/1     0 0   0   dev12       6
10.15.72.253    10.15.72.99     6    3488  5001   1/1     0 0   0   dev5        6    C
10.15.72.99     10.15.72.253    6    5001  3488   1/2     0 0   0   dev5        6
10.15.72.99     10.15.72.253    6    36092 5001   1/2     0 0   0   dev12       6    C
```

The following is a sample output for a per-user configuration:

```
(host) #show datapath session ap-name rap5-2 table 10.15.72.99

Datapath Session Table Entries
------------------------------
Flags: F - fast age, S - src NAT, N - dest NAT
D - deny, R - redirect, Y - no syn
H - high prio, P - set prio, T - set ToS
C - client, M - mirror, V - VOIP
Q - Real-Time Quality analysis
I - Deep inspect, U - Locally destined
E - Media Deep Inspect, G - media signal
RAP Flags: 1 - Class 1, 2 - Class 2, 3 - Class 3
Source IP     Destination IP  Prot SPort DPort  Cntr Prio ToS Age Destination TAge Flags
--------------  --------------  ---- ----- -----  ---- ---- --- --- ----------- ---- -----
10.15.72.253    10.15.72.99     6    3489  5001   1/3     0 0   0   dev5        37   FC
10.15.72.99     10.15.72.253    6    5001  3489   1/4     0 0   0   dev5        37   F
10.15.72.99     10.15.72.253    6    36096 5001   1/4     0 0   0   dev12       37   C
10.15.72.253    10.15.72.99     6    5001  36096  1/3     0 0   0   dev12       37
```

Virtual Intranet Access (VIA) is part of the Dell W-Series remote networks solution targeted for teleworkers and mobile users. VIA detects the users network environment (trusted and un-trusted) and automatically connects the user to their enterprise network. Trusted networks typically refers to a protected office network that allows users to directly access corporate intranet. Un-trusted networks are public Wi-Fi hotspots like airports, cafes, or home network. The VIA solution comes in two parts—VIA connection manager and the controller configuration.

- VIA connection manager—Teleworkers and mobile users can easily install a light weight application on their Microsoft Windows or Apple MacBook computers to connect to their enterprise network from remote locations (see Understanding VIA Connection Manager on page 679).
- Controller configuration—To set up virtual intranet access for remote users, you must configure your controller to include setting up user roles, authentication, and connection profiles. You can use either WebUI or CLI to configure your controller (see Configuring the VIA Controller on page 682).



VIA requires the PEFV license and is supported on the W-600 Series, W-3000 Series, W-6000M3, and W-7200 controller.

Topics in this chapter include:

- Understanding VIA Connection Manager on page 679
- Downloading VIA on page 696
- Configuring the VIA Controller on page 682

Understanding VIA Connection Manager

If a user is connected from a remote location that is outside of the enterprise network, VIA automatically detects the environment as un-trusted and creates a secure IPSec connection between the user and the enterprise network. When the user moves into the trusted network, VIA detects the network type and moves to idle state.

How it Works

VIA provides a seamless connectivity experience to users when accessing an enterprise network resource from an un-trusted or trusted network environment. You can securely connect to your enterprise network from an un-trusted network environment. By default VIA will auto-launch at system start and establish a remote connection. The following table explains the typical behavior:



The sequence of events described in Table 127 does not necessarily mean that the events always happen in the order shown in the table.

Table 127: *VIA Connectivity Behavior*

User action / environment	VIA's behavior
The client / user moves from a trusted to un-trusted environment. Example: From office to a public hot-spot.	Auto-launches and establishes connection to remote network.
The client moves from an un-trusted to a trusted environment.	Auto-launch and stay idle. VIA does not establish remote connection. You can, however, manually connect to a network by selecting an appropriate connection profile from the *Settings* tab.
While in an un-trusted environment, user disconnects the remote connection.	Disconnects gracefully.
User moves to a trusted environment.	Stays idle and does not connect.
User moves to an un-trusted environment	Stays idle and does not connect. This usually happens, if the user has in a previous occasion disconnected a secure connection by clicking the **Disconnect** button in VIA. Users can manually connect by one of the following methods: 1. Right click on the VIA icon in the system tray and select the **Restore** option and then select the **Connect** option to connect using the default connection profile. 2. Right click on the VIA icon in the system tray and select the **Connect** option.
User clicks the **Reconnect** button.	Establishes remote connection.
In an un-trusted environment, user restarts the system.	Auto-launches and establishes remote connection.
In an un-trusted environment, user shuts down the system. Moves to a trusted environment and restarts system.	Auto-launches and stays idle.

Installing the VIA Connection Manager

Users can download VIA from a URL provided to them by their IT department and install it on their computers.

On Microsoft Windows Computers

1. Download the installer (*ansetup.msi or ansetup64.msi*) from the URL provided by the IT department.
2. Double click the installer file and follow the default prompts.
3. After the installation is complete, the user will be prompted to enter the following:
   a. Remote server URL—This should be provided by the IT department. The administrator can also provision the URL on the controller. In such cases, the user is required to specify only the username and password.
   b. Username—The users domain user name.
   c. Password—The users domain password.
4. Click the **Connect** button to initiate a secure VIA connection. VIA will minimized to system tray after establishing the secure connection.

On Apple MacBooks

1. Download the installer (*ansetup.dmg*) from the URL provided by the IT department.

2. Double click the installer file and follow the default prompts.
3. After the installation is complete, the user will be prompted to enter the following:
   a. Remote server URL—This should be provided by the IT department.
   b. Username—The users domain user name.
   c. Password—The users domain password.
4. Go to **System Preferences** > **Other** > select **VIA** to view VIA connection details.
5. Go to **System Preferences** > **Network**, in the list of network connections select **VIA** to modify login details and remote server address.

Upgrade Workflow

VIA checks for upgrade requirements during the login phase. There are two types of upgrade process: Minimal Upgrade and Complete Upgrade.

Minimal Upgrade

This type of upgrade is initiated for bug fixes and some minor enhancements which requires only some components of the client to be upgraded. When a VPN session is active the upgrade binary is downloaded by VIA from the controller. After the active VIA connection is terminated, the upgrade process is started and the client is upgraded. This type of upgrade does not require a system reboot.

Complete Upgrade

This requires an upgrade to VIA and its underlying network drivers. This type of upgrade requires a system reboot. VIA downloads the upgrade binary from the controller and displays a message about upgrade process after the connection is terminated for that upgrade. The user can choose to proceed or cancel the upgrade process. If the user chooses to upgrade, a system reboot is automatically executed. If the user cancels the upgrade, VIA will prompt the user for an upgrade every time the user terminates a VIA session.



See Downloading VIA on page 696 for information about using the desktop application.

VIA Compatibility Matrix

The following table shows the compatibility of different versions of VIA with ArubaOS.

Table 128: *VIA Compatibility Matrix*

ArubaOS Version	Microsoft Windows (32/64 bit) XP (32-bit only) Vista Windows 7/8/8.1	Mac OS 10.7, 10.8, 10.9	Apple iOS 4.x, 5.x, 6.x, 7.x	Android 4.0, 4.1, 4.2, 4.3, 4.4	Linux (32/64 bit) Ubuntu 12.04 RHEL 6 CentOS 6
ArubaOS 5.0.x	2.0.x.x, 2.1.x.x	—	—	—	—
ArubaOS 6.0.x	2.0.x.x, 2.1.x.x	2.0.x	2.0.x, 2.1.x	2.1.x.x, 2.2.x.x	2.0.x
ArubaOS 6.1.x	2.0.x.x, 2.1.x.x	2.0.x	2.0.x, 2.1.x	2.1.x.x, 2.2.x.x	2.0.x
ArubaOS 6.2.x	2.0.x.x, 2.1.x.x	2.0.x	2.0.x, 2.1.x	2.1.x.x,	2.0.x

Table 128: *VIA Compatibility Matrix*

ArubaOS Version	Microsoft Windows (32/64 bit) XP (32-bit only) Vista Windows 7/8/8.1	Mac OS 10.7, 10.8, 10.9	Apple iOS 4.x, 5.x, 6.x, 7.x	Android 4.0, 4.1, 4.2, 4.3, 4.4	Linux (32/64 bit) Ubuntu 12.04 RHEL 6 CentOS 6
				2.2.x.x	
ArubaOS 6.3.x	2.0.x.x, 2.1.x.x	2.0.x	2.0.x, 2.1.x	2.1.x.x, 2.2.x.x	2.0.x
ArubaOS 6.4.x	2.0.x.x, 2.1.x.x	2.0.x	2.0.x, 2.1.x	2.1.x.x, 2.2.x.x	2.0.x

Configuring the VIA Controller

VIA configuration requires that you first configure VPN settings and then configure VIA settings. See Virtual Private Networks on page 337 for information on configuring VPN settings on your controller.

Before you Begin

The following ports must be enabled before configuring the VIA controller.

- **TCP 443**—During the initializing phase, VIA uses HTTPS connections to perform trusted network and captive portal checks against the controller. It is mandatory that you enable port 443 on your network to allow VIA to perform these checks.
- **UDP 4500**—Required for IPSec transport
- **UDP 500**—Required for VIA 1.0 on Mac OS

Supported Authentication Mechanisms

VIA 1.x and VIA 2.x support different authentication mechanisms:

Authentication mechanisms supported in VIA 1.x

Authentication is performed using IKEv1 only. Phase 0 authentication, which authenticates the VPN client, can be performed using either a pre-shared key or an X.509 certificate (the X.509 certificate must appear in the operating system's "user" certificate store.). If certificates are used for IKE phase 0 authentication, it must be followed by username/password authentication.

The second authentication phase is performed using xAuth, which requires a username and password. The username and password is authenticated against the controller's internal database, a RADIUS server, or an LDAP server. If a RADIUS server is used, it must support the PAP or MSCHAPv2 protocol. By default, PAP protocol is enabled for RADIUS authentication.

Support for two-factor authentication such as token cards is provided in VIA 1.x. Token product like RSA tokens and other token cards are also supported. This includes support for new-pin and next-pin.

Authentication mechanisms supported in VIA 2.x

In addition to the authentication methods supported by VIA 1.x, VIA 2.x adds support for IKEv2. IKEv2 is an updated version that is faster and supports a wider variety of authentication mechanisms. IKEv2 does not have two phases of authentication, only a single phase. VIA supports the following with IKEv2:

- Username/password
- X.509 certificate. Controllers running ArubaOS 6.1 or greater support OCSP for the purpose of validating that a certificate has not been revoked.
- EAP (Extensible Authentication Protocol) including EAP-TLS and EAP-MSCHAPv2.

Other authentication methods:

- Certificates based authentication.
- Smart cards that support a Smart Card Cryptographic Provider (SCCP) API within the operating system. VIA will look for an X.509 certificate in the operating system's certificate store. A smart card supporting a SCCP will cause the certificate embedded within the smart card to automatically appear in the operating system's certificate store.

Suite B Cryptography Support

Suite B is a new set of cryptographic algorithms that are approved by the US Government for use in classified communication. Suite B provides the highest levels of security available today in public, commercial algorithms. Specifically, VIA provides support for:

- RFC 4869—Suite B Cryptographic Suites for IPsec
- AES-GCM 128/256 for bulk data transfer
- ECDSA for digital signatures, including support for X.509v3 certificates using ECDSA keys with p256/p384 curves
- ECDH for key agreement using p256/p384 curves
- SHA-256 and SHA-384 for message digests



Suite B support requires a controller running Dell Networking W-Series ArubaOS 6.4.x or greater with the Advanced Cryptography License installed. See Software Licenses on page 130 for more information on licenses.

802.11 Suite-B

The bSec protocol is a pre-standard protocol that has been proposed to the IEEE 802.11 committee as an alternative to 802.11i. The main difference between bSec and standard 802.11i is that bSec implements Suite B algorithms wherever possible. Notably, AES-CCM is replaced by AES-GCM, and the Key Derivation Function (KDF) of 802.11i is upgraded to support SHA-256 and SHA-384. In order to provide interoperability with standard Wi-Fi software drivers, bSec is implemented as a shim layer between standard 802.11 Wi-Fi and a Layer 3 protocol such as IP. A controller configured to advertise a bSec SSID will advertise an open network, however only bSec frames will be permitted on the network.

The bSec protocol requires that you use VIA 2.1. or greater on the client device.

Configuring VIA Settings

The following steps are required to configure your controller for VIA. These steps are described in detail in the subsections that follow.

1. Enable VPN Server Module—ArubaOS allows you to connect to the VIA controller using the default user roles. However, to configure and assign specific user roles you must install the Policy Enforcement Firewall Virtual Private Network (PEFV) license. For details, see Enable VPN Server Module on page 684.
2. Create VIA User Roles—VIA user roles contain access control policies for users connecting to your network using VIA. You can configure different VIA roles or use the default VIA role—`default-via-role`.For details, see Create VIA User Roles on page 684.
3. Create VIA Authentication Profile—A VIA authentication profile contains a server group for authenticating VIA users. The server group contains the list of authentication servers and server rules to derive user roles

based on the user authentication. You can configure multiple VIA authentication profiles and / or use the default VIA authentication profile created with *Internal* server group. For details, see Create VIA Authentication Profile on page 684.

4. Create VIA Connection Profile— A VIA connection profile contains settings required by VIA to establish a secure connection to the controller. You can configure multiple VIA connection profiles. A VIA connection profile is always associated to a user role and all users belonging to that role will use the configured settings. If you do not assign a VIA connection profile to a user role, the default connection profile is used.For details, see Create VIA Connection Profile on page 686.
5. Configure VIA Web Authentication—A VIA web authentication profile contains an ordered list of VIA authentication profiles. The web authentication profile is used by end users to login to the VIA download page (*https://<server-IP-address>/via*) for downloading the VIA client. Only one VIA web authentication profile is available. If more than one VIA authentication profile is configured, users can view this list and select one during the client login. For details, see Configure VIA Web Authentication on page 690.
6. Associate VIA Connection Profile to User Role—A VIA connection profile has to be associated to a user role. Users will login by authenticating against the server group specified in the VIA authentication profile and are put into that user role. The VIA configuration settings are derived from the VIA connection profile attached to that user role. Default connection profile is used. For details, see Associate VIA Connection Profile to User Role on page 691.
7. Configure VIA Client WLAN Profiles—You can push WLAN profiles to end-user computers that use the Microsoft Windows Wireless Zero Config (WZC) service to configure and maintain their wireless networks. After the WLAN profiles are pushed to end-user computers, they are automatically displayed as an ordered list in the preferred networks. The VIA client WLAN profiles provisioned on the client can be selected from the VIA connection profile described in Step 6. For details, see Configure VIA Client WLAN Profiles on page 691.
8. Rebranding VIA and Downloading the Installer—You can use a custom logo on the VIA client and on the VIA download web page. For details, see Rebranding VIA and Downloading the Installer on page 694.
9. Download VIA Installer and Version File

Using the WebUI to Configure VIA

The following steps illustrate configuring your controller for VIA using the WebUI.

Enable VPN Server Module

You must install the PEFV license to configure and assign user roles. See Software Licenses on page 130 for licensing requirements.

To install a license:

1. Navigate to **Configuration** > **Network** > **Controller** and select the **Licenses** tab on the right hand side.
2. Paste the license key in the **Add New License key** text box and click the **Add** button.

Create VIA User Roles

To create VIA users roles:

1. Navigate to **Configuration** > **Security** > **Access Control** > **User Roles**.
2. Click **Add** to create new policies. Click **Done** after creating the user role and apply to save it to the configuration.

Create VIA Authentication Profile

This following steps illustrate the procedure to create an authentication profile to authenticate users against a server group.

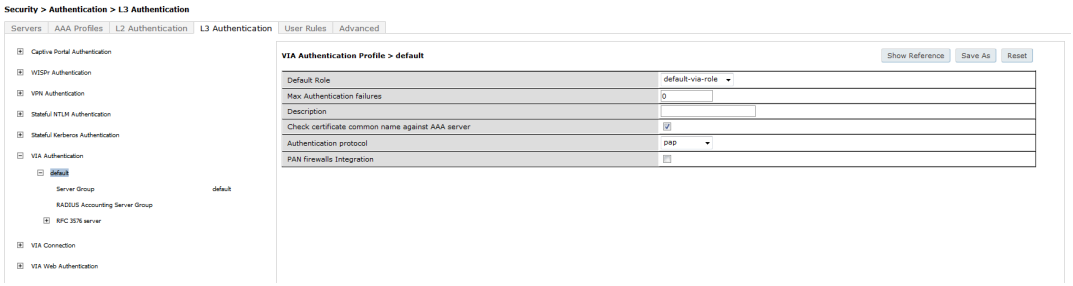
1. Navigate to **Configuration** > **Security** > **Authentication** > **L3 Authentication**.
2. Under the *Profiles* section, expand the **VIA Authentication** option. You can configure the following parameters for the authentication profile:

Table 129: *VIA - Authentication Profile Parameters*

Parameter	Description
Default Role	The role that will be assigned to the authenticated users.
Max Authentication Failures	Specifies the maximum authentication failures allowed. The default is 0 (zero).
Description	A user friendly name or description for the authentication profile.
Check certificate common name against AAA server	If you use client certificates for user authentication, enable this option to verify that the certificate's common name exists in the server.
Authentication protocol	PAP and MSCHAPv2 protocols are supported to authenticate VIA users. Default: PAP
PAN firewalls Integration	Requires IP mapping at Palo Alto Networks firewalls.

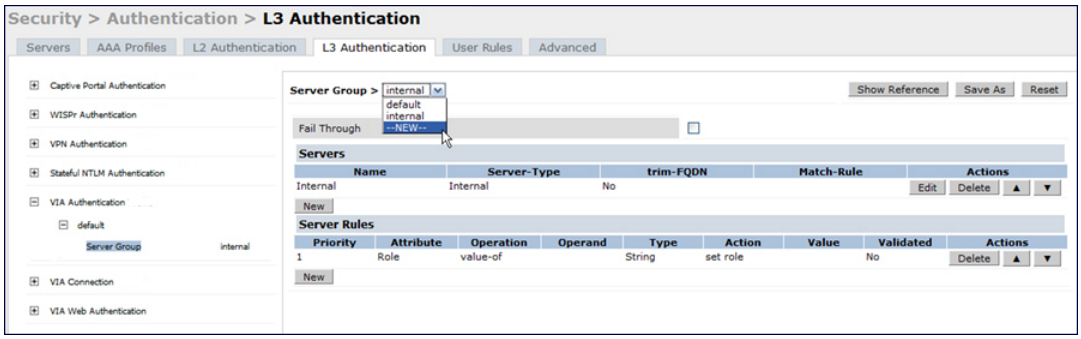
3. To create a new authentication profile:
   a. Enter a name for the new authentication profile under the *VIA Authentication* section and click the **Add** button.
   b. Expand the **VIA Authentication** option and select the new profile name.
4. To modify an authentication profile, select the profile name to configure the default role
   The following screenshot uses the default authentication profile.

Figure 89 *VIA - Associate User Role to VIA Authentication Profile*



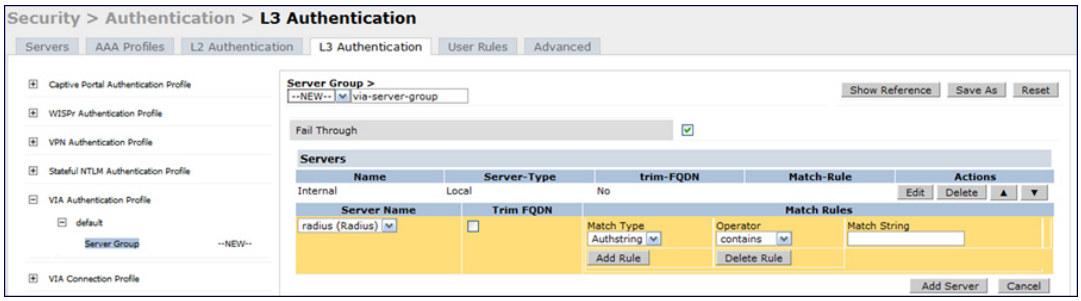
5. To use a different server group, Click *Server Group* under VIA Authentication Profile and select **New** to create a new server group.

Figure 90 *VIA - Creating a new server group for VIA authentication profile*



6. Enter a name for the server group.

Figure 91 *VIA - Enter a name for the server group*

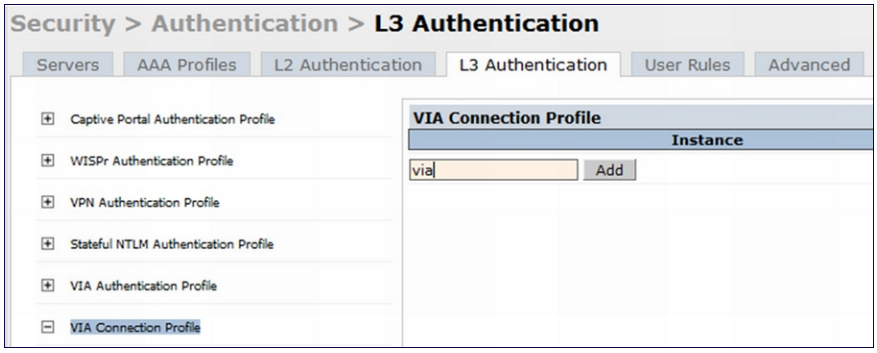


Create VIA Connection Profile

To create VIA connection profile:

1. Navigate to **Configuration** > **Security** > **Authentication** > **L3 Authentication** tab. Click the VIA *Connection Profile* option and enter a name for the connection profile.

Figure 92 *VIA - Create VIA Connection Profile*



2. Click on the new VIA connection profile to configure the connection settings. VIA Connection profile settings are divided into two tabs, **Basic** and **Advanced**. The **Basic** tab displays only those configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. If you change a setting on one tab then click and display the other tab without saving your configuration, that setting will revert to its previous value.
3. Click **Apply** to save your changes.

Table 130: *VIA - Connection Profile Options*

Configuration Option	Description
Basic VIA Connection Profile Settings	
VIA Servers	Enter the following information about the VIA controller. • *Controller Hostname/IP Address*: This is the public IP address or the DNS hostname of the VIA controller. Users will connect to remote server using this IP address or the hostname. • Controller Internal IP Address: This is the IP address of any of the VLAN interface IP addresses belongs to this controller. • Controller Description: This is a human-readable description of the controller. Click the **Add** button after you have entered all the details. If you have more than one VIA controller you order them by clicking the *Up* and *Down* arrows. To delete a controller from your list, select a controller and click the **Delete** button.
Client Auto-Login	Enable or disable VIA client to auto login and establish a secure connection to the controller. Default: Enabled
VIA tunneled networks	A list of network destination (IP address and netmask) that the VIA client will tunnel through the controller. All other network destinations will be reachable directly by the VIA client. • Enter an IP address & network mask and click the **Add** button to add to the tunneled networks list. • To delete a network entry, select the IP address and click the **Delete** button.
Enable split-tunneling	Enable or disable split tunneling. • If enabled, all traffic to the VIA tunneled networks (Step 3 in this table) will go through the controller and the rest is just bridged directly on the client. • If disabled, all traffic will flow through the controller. Default: off
Allow client-side logging	Enable or disable client side logging. If enabled, VIA client will collect logs that can be sent to the support email-address for troubleshooting. Default: Enabled
Enable IKEv2	Select this option to enable or disable the use of IKEv2 policies for VIA.
Use Suite B Cryptography	Select this option to use Suite B cryptography methods. You must install the Advanced Cryptography license to use the Suite B cryptography. See Working with Licenses on page 131 for more information.
IKEv2 Authentication method	List of all IKEv2 authentication methods.
VIA Client DNS Suffix List	The DNS suffix list (comma separated) that has be set on the client once the VPN connection is established. Default: None.
VIA Support E-mail Address	The support e-mail address to which VIA users will send client logs. Default: None.
Advanced VIA Connection Profile Settings	
VIA Servers	Enter the following information about the VIA controller.

Configuration Option	Description
	• **Hostname/IP Address**: This is the public IP address or the DNS hostname of your VIA Server / controller. Users will connect to this remote server using the IP address or the hostname. • **Internal IP Address**: This is the IP address of any of the VLAN interface IP addresses belonging to this VIA server. • **Description**: This is a human-readable description of the VIA server. Click the **Add** button after you have entered all the details. If you have more than one VIA controller you re-order them by clicking the Up and Down arrows. To delete a VIA server from your list, select a server and click **Delete**.
Client Auto-Login	Select this checkbox to allow a VIA client to automatically log in and establish a secure connection to the controller. Default: Enabled
VIA Authentication Profiles to provision	This is the list of VIA authentication profiles that will be displayed to users in the VIA client. See Create VIA Connection Profile on page 686 • Select an authentication profile and click the Add button to add to the authentication profiles list. • You can change the order of the list by clicking the *Up* and *Down* arrows. • To delete an authentication profile, select a profile name and click the **Delete** button.
Allow client to auto-upgrade	Enable or disable VIA client to automatically upgrade when an updated version of the client is available on the controller. Default: Enabled
VIA tunneled networks	A list of network destination (IP address and netmask) that the VIA client will tunnel through the controller. All other network destinations will be reachable directly by the VIA client. Enter an IP address and network mask, then click **Add** button to add them to the tunneled networks list. To delete a network entry, select the IP address and click **Delete**.
Enable split-tunneling	Enable or disable split tunneling. If enabled, all traffic to the VIA tunneled networks) will go through the controller and the rest is just bridged directly on the client. If disabled, all traffic will flow through the controller. Default: off
VIA Client WLAN profiles	A list of VIA client WLAN profiles that needs to be pushed to the client machines that use Windows Zero Config (WZC) to configure or manage their wireless networks. • Select a WLAN profile and click the **Add** button to add to the client WLAN profiles list. • To delete an entry, select the profile name and click the **Delete** button. See Configure VIA Client WLAN Profiles on page 691 for more information.
VIA IKE V2 Policy	List of available IKEv2 policies.
VIA IKE Policy	List of IKE policies that the VIA Client has to use to connect to the controller. These IKE policies are configured under **Configuration** > **Advanced Services** > **VPN Services** > **IPSEC** > **IKE Policies**.
Use Windows Credentials	Enable or disable the use of the Windows credentials to login to VIA. If enabled, the SSO (Single Sign-on) feature can be utilized by remote users to connect to internal resources. Default: Enabled

Configuration Option	Description
VIA IPSec V2 Crypto Map	List of all IPSec V2 that the VIA client uses to connect to the controller.
VIA IPSec Crypto Map	List of IPSec Crypto Map that the VIA client uses to connect to the controller. These IPSec Crypto Maps are configured in CLI using the `crypto-local ipsec-map <ipsec-map-name>` command.
VIA Client Network Mask	The network mask that has to be set on the client after the VPN connection is established. Default: 255.255.255.255
Enable Supplicant	If enabled, VIA starts in bSec mode using L2 suite-b cryptography. This option is disabled by default.
Enable FIPS Module	Enable the VIA (Federal Information Processing Standard) FIPS module so VIA checks for FIPS compliance during startup. This option is disabled by default.
Auto-Launch Supplicant	Select this option to automatically connect to a configured WLAN network.
Lockdown all Settings	If enabled, all user options on the VIA client are disabled.
Domain Suffix in VIA Authentication	Enables a domain suffix on VIA Authentication, so client credentials are sent as domainname\username instead of just username..
Enable Controllers Load Balance	Enable this option to allow the VIA client to failover to the next available selected randomly from the list as configured in the VIA Servers option. If disabled, VIA will failover to the next in the sequence of ordered list of VIA Servers.
Enable Domain Pre-connect	Enable this option to allow users with lost or expired passwords to establish a VIA connection to corporate network. This option authenticates the user's device and establishes a VIA connection that allows users to reset credentials and continue with corporate access.
VIA Banner Message Reappearance Timeout (minutes)	The maximum time (minutes) allowed before the VIA login banner reappears. Default: 1440 min
VIA Client Network Mask	VIA client network mask, in dotted decimal format.
Validate Server Certificate	Enable or disable VIA from validating the server certificate presented by the controller. Default: Enabled
VIA max session timeout	The maximum time (minutes) allowed before the VIA session is disconnected. Default: 1440 min
VIA Logon Script	Specify the name of the logon script that must be executed after VIA establishes a secure connection. The logon script must reside in the client computer.
VIA Logoff Script	Specify the name of the log-off script that must be executed after the VIA connection is disconnected. The logoff script must reside in the client computer.

Configuration Option	Description
Maximum reconnection attempts	The maximum number of re-connection attempts by the VIA client due to authentication failures. Default: 3
VIA external download URL	End users will use this URL to download VIA on their computers.
Allow user to disconnect VIA	Enable or disable users to disconnect their VIA sessions. Default: on
Comma separated list of HTTP ports to be inspected (apart from default port 80)	Traffic from the specified ports will be verified by the content security service provider.
Keep VIA window minimized	Enable this option to minimize the VIA client to system tray during the connection phase. Applicable to VIA client installed in computers running Microsoft Windows operating system.
Block traffic until VPN tunnel is up	If enabled, this feature will block network access until the VIA VPN connection is established.
Block traffic rules	Specify a hostname or IP address and network mask to define a whitelist of users to which the **Block traffic until VPN tunnel is up** setting will not apply.
User idle timeout	Select the **Enable** checkbox to configure user idle timeout value for this profile. Specify the idle timeout value for the client in seconds. Valid range is 30-15300 in multiples of 30 seconds. Enabling this option overrides the global settings configured in the AAA timers. If this is disabled, the global settings are used.

Configure VIA Web Authentication

To configure VIA web authentication profile:

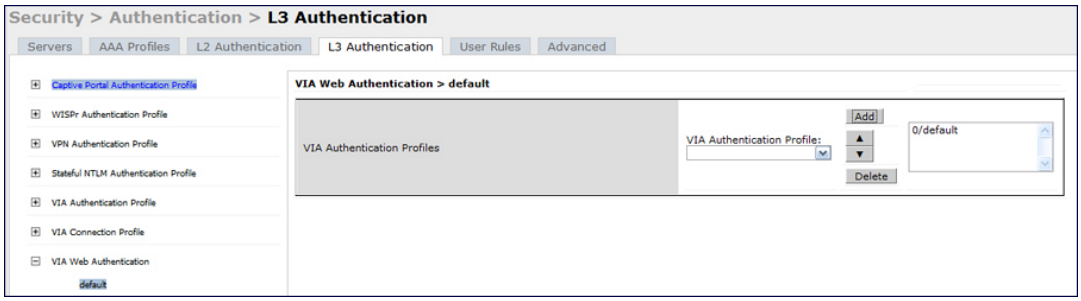
1. Navigate to **Configuration** > **Security** > **Authentication** > **L3 Authentication** tab.
2. Expand VIA Web Authentication and click on *default* profile.



You can have only one profile (`default`) for VIA web authentication.

3. Select a profile from **VIA Authentication Profile** drop-down list box and click the **Add** button.
   - To re-order profiles, click the *Up* and *Down* button.
   - To delete a profile, select a profile and click the **Delete** button.
4. If a profile is not selected, the *default* VIA authentication profile is used.

Figure 93 *VIA - Select VIA Authentication Profile*

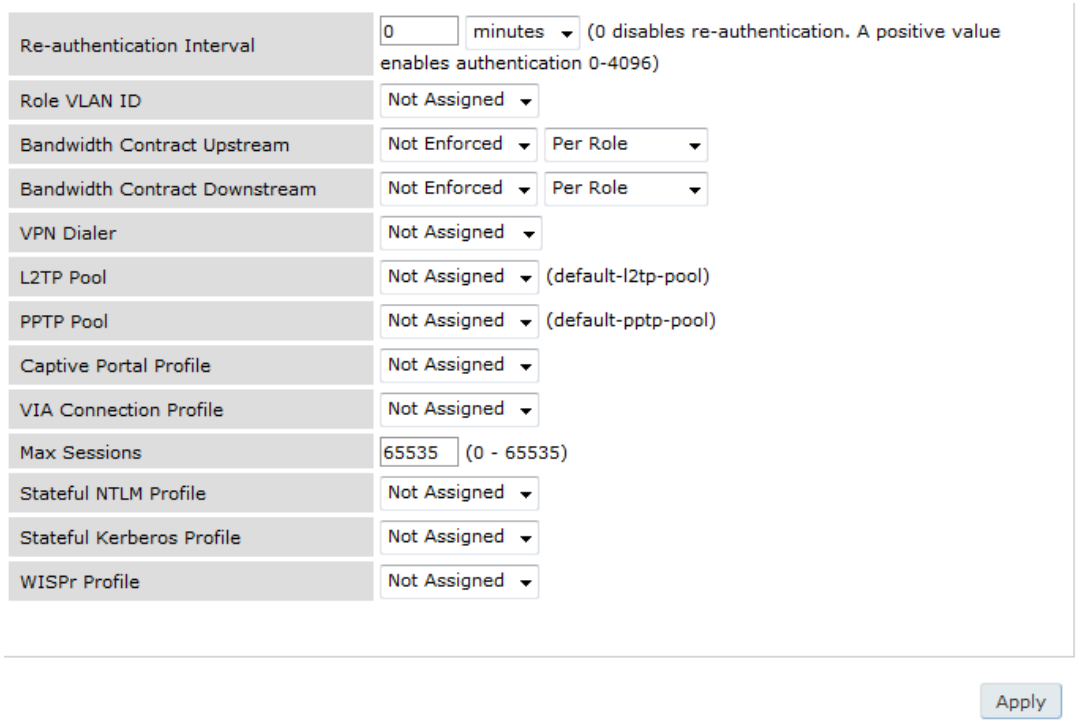


Associate VIA Connection Profile to User Role

To associate a VIA connection profile to a user role:

1. Navigate to **Configuration** > **Security** > **Access Control** > **User Roles** tab.
2. Select the VIA user role (See Create VIA User Roles on page 684) and click the **Edit** button.
3. In the *Edit Role* page, navigate to VIA Connection Profile and select the connection profile from the drop-down list box.
4. Click the **Apply** button to save the changes to the configuration.

Figure 94 *VIA - Associate VIA Connection Profile to User Role*

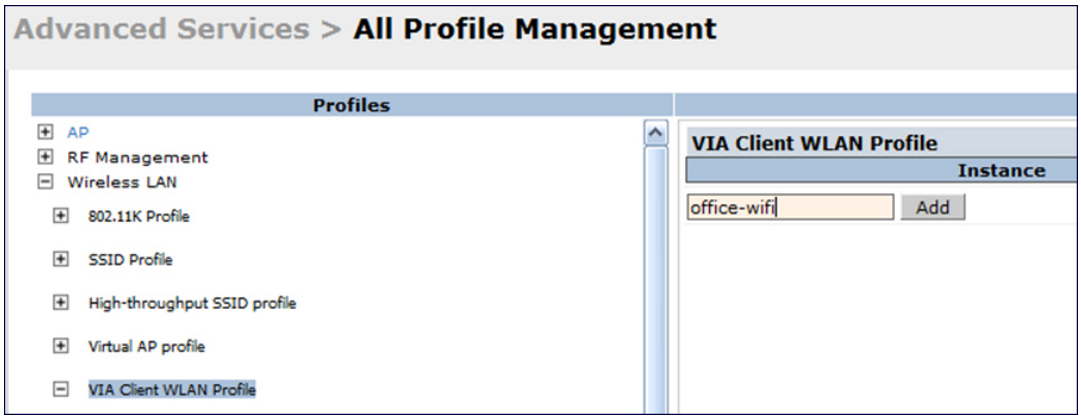


Configure VIA Client WLAN Profiles

To configure a VIA client WLAN profile:

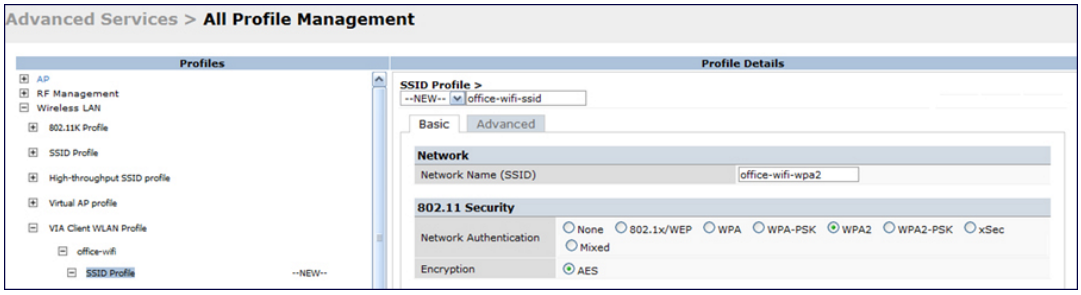
1. Navigate to **Configuration** > **Advanced Services** > **All Profiles**.
2. Expand Wireless LAN and select **VIA Client WLAN** .
3. In the Profile Details, enter a name for the WLAN profile and click the **Add** button.

Figure 95 *VIA - Create VIA Client WLAN Profile*



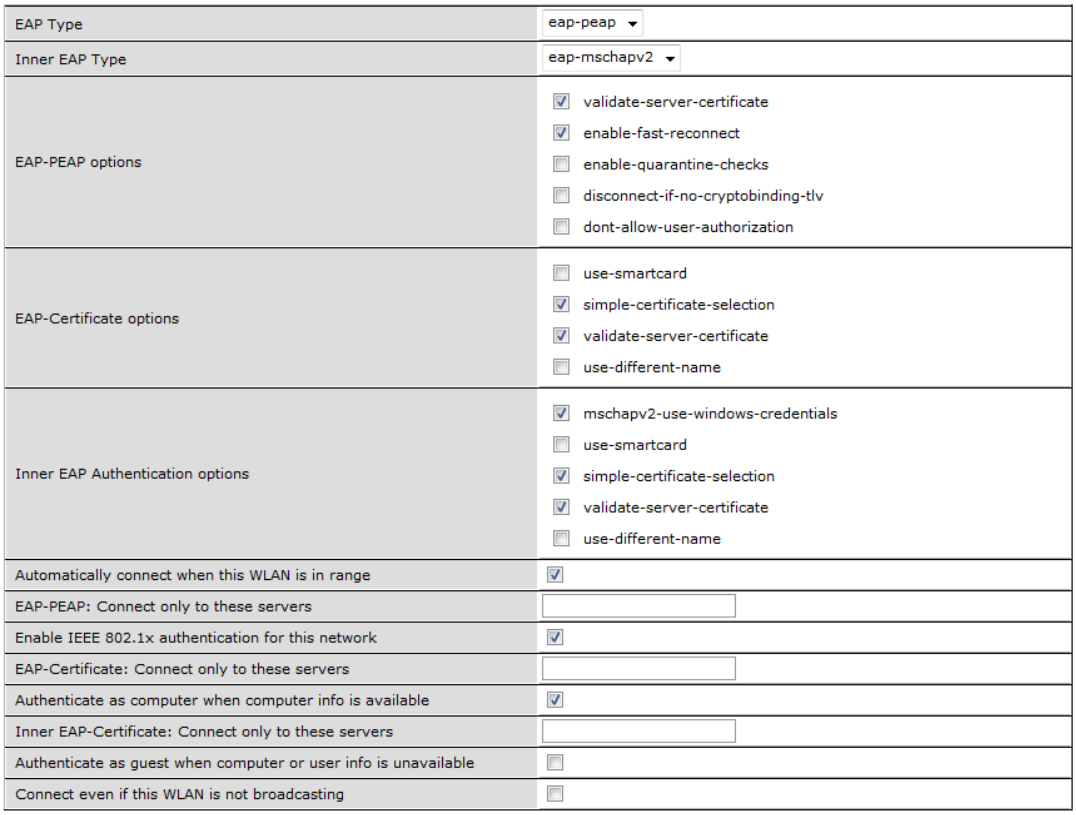
4. Expand the new WLAN profile and click SSID Profile. In the profile details page, select **New** from the SSDI Profile drop-down box and enter a name for the SSID profile.
5. In the Basic tab, enter the network name (SSID) and select 802.11 security settings. Click the **Apply** button to continue.

Figure 96 *VIA - Configure the SSID Profile*



6. You can now configure the SSID profile by selecting the SSID profile under VIA Client WLAN Profile option.

Figure 97 *VIA - Configure VIA Client WLAN Profile*



The VIA client WLAN profile are similar to the authentication settings used to set up a wireless network in Microsoft Windows. The following table shows the Microsoft Windows equivalent settings:

Table 131: *Configure VIA client WLAN profile*

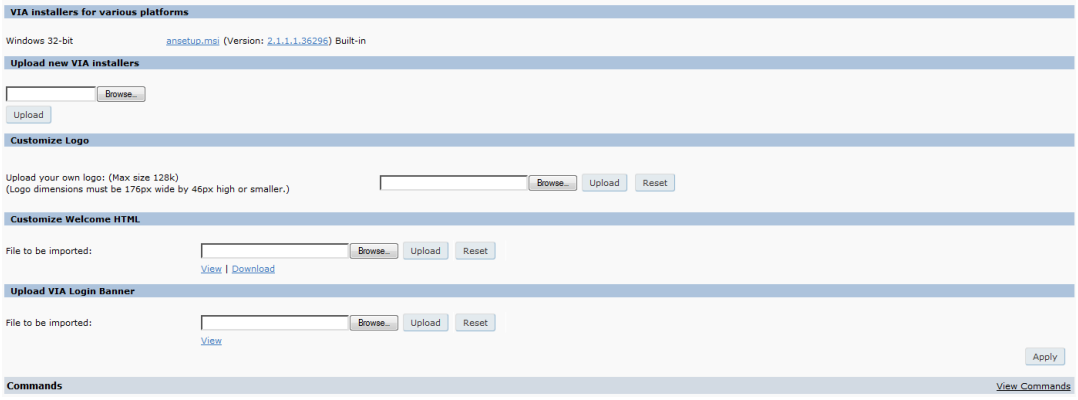
Option	Description
EAP-PEAP options	Select the following options, if the EAP type is PEAP (Protected EAP): • validate-server-certificate: Select this option to validate server certificates. • enable-fast-reconnect: Select this option to allow fast reconnect. • enable-quarantine-checks: Select this option to perform quarantine checks. • disconnect-if-no-cryptobinding-tlv: Select this option to disconnect if server does not present cryptobinding TLV. • dont-allow-user-authorization: Select this to disable prompts to user for authorizing new servers or trusted certification authorities.
EAP Type	Select an EAP type used by client to connect to wireless network. Default: EAP-PEAP
EAP-Certificate Options	If you select EAP type as certificate, you can select one of the following options: • mschapv2-use-windows-credentials • use-smartcard • simple-certificate-selection • use-different-name • validate-server-certificate
Inner EAP Type	Select the inner EAP type. Default: EAP-MSCHAPv2

Option	Description
Inner EAP Authentication options:	• mschapv2-use-windows-credentials: Automatically use the Windows logon name and password (and domain if any) • use-smartcard: Use a smart card • simple-certificate-selection: Use a certificate on the users computer or use a simple certificate selection method (recommended) • validate-server-certificate: Validate the server certificate • use-different-name: Use a different user name for the connection (and not the CN on the certificate)
Automatically connect when this WLAN is in range	Select this option if you want VIA client to connect when this network (SSID) is available.
EAP-PEAP: Connect only to these servers	Comma separated list of servers.
Enable IEEE 802.1X authentication for this network	Select this option to enable 802.1X authentication for this network. Default: Enabled.
EAP-Certificate: Connect only to these certificates	Comma separated list of servers.
Inner EAP-Certificate: Connect only to these servers	Comma separated list of servers.
Connect even if this WLAN is not broadcasting	Default: Disabled

Rebranding VIA and Downloading the Installer

You can re-brand the VIA client and the VIA download page with your custom logo and HTML page.

Figure 98 *VIA - Customize VIA logo, Landing Page, and download VIA Installer*



Download VIA Installer and Version File

To download the VIA installer and version file:

1. Navigate to the **Configuration > Advanced Services > VPN Services > VIA** tab.
2. Under VIA installers for various platforms section, click **`ansetup.msi`** to download the installation file.

Customize VIA Logo

To use a custom logo on the VIA download page and the VIA client:

1. Navigate to the**Configuration > Advanced Services > VPN Services > VIA** tab.
2. Under *Customize Logo* section, browse and select a logo from your computer. Click the **Upload** button to upload the image to the controller.
   - To use the default Dell logo, click the **Reset** button.

Customize the Landing Page for Web-based Login

To use a custom landing page for VIA web login:

1. Navigate to the **Configuration > Advanced Services > VPN Services > VIA** tab.
2. Under Customize Welcome HTML section, browse and select the HTML file from your computer. Click the **Upload** button to upload the image to the controller.
3. The following variables are used in the custom HTML file:

   All variables in the custom HTML file have the following notation
   - `<% user %>`: this will display the username.
   - `<% ip %>`: this will display the IP address of the user.
   - `<% role %>`: this will be display the user role.
   - `<% logo %>`: this is the custom logo (Example: `<img src="<% logo %>">`)
   - `<% logout %>`: the logout link (Example: `<a href="<% logout %>">VIA Web Logout</a>`)
   - `<% download %>`: the installer download link (Example: `<a href="<% download %>">Click here to download VIA</a>`)

   To use the default welcome page, click the **Reset** button.
4. Click the **Apply** button to continue.

Using the CLI to Configure VIA

The following steps illustrate configuring VIA Using the CLI. Install your Policy Enforcement Firewall Virtual Private Network (PEFV) license key. For detailed information on the VIA command line options, see the *Dell Networking W-Series ArubaOS 6.4.x Command Line Interface Guide*.

```
(host) (config)# license add <key>
```

Create VIA roles

```
(host) (config) #user-role example-via-role
(host) (config-role) #access-list session "allowall" position 1
(host) (config-role) #ipv6 session-acl "v6-allowall" position 2
```

Create VIA authentication profiles

```
(host) (config) #aaa server-group "via-server-group"
(host) (Server Group "via-server-group") #auth-server "Internal" position 1
(host) (Server Group "via-server-group") #aaa authentication via auth-profile default
(host) (VIA Authentication Profile "default") #default-role example-via-role
(host) (VIA Authentication Profile "default") #desc "Default VIA Authentication Profile"
(host) (VIA Authentication Profile "default") #server-group "via-server-group"
```

Create VIA connection profiles

```
(host) (config) #aaa authentication via connection-profile "via"
(host) (VIA Connection Profile "via") #server addr 202.100.10.100 internal-ip 10.11.12.13 desc
"VIA Primary" position 0
(host) (VIA Connection Profile "via") #auth-profile "default" position 0
(host) (VIA Connection Profile "via") #tunnel address 10.0.0.0 netmask 255.255.255.0
```

```
(host) (VIA Connection Profile "via") #split-tunneling
(host) (VIA Connection Profile "via") #windows-credentials
(host) (VIA Connection Profile "via") #client-netmask 255.0.0.0
(host) (VIA Connection Profile "via") #dns-suffix-list example.com
(host) (VIA Connection Profile "via") #support-email via-support@example.com
```

Enter the following command after you create the client WLAN profile. See Configure VIA Client WLAN Profiles on page 691

```
(host) (VIA Connection Profile "via") #client-wlan-profile "via_corporate_wpa2" position 0
```

Configure VIA web authentication

```
(host) (config) #aaa authentication via web-auth default
(host) (VIA Web Authentication "default") #auth-profile default position 0
```



You can have only one profile (*default*) for VIA web authentication.

Associate VIA connection profile to user role

```
(host) (config) #user-role "example-via-role"
(host) (config-role) #via "via"
```

Configure VIA client WLAN profiles

```
(host) (config) #wlan ssid-profile "via_corporate_wpa2"
(host) (SSID Profile "via_corporate_wpa2") #essid corporate_wpa2
(host) (SSID Profile "via_corporate_wpa2") #opmode wpa2-aes
(host) (SSID Profile "via_corporate_wpa2") #wlan client-wlan-profile "via_corporate_wpa2"
(host) (VIA Client WLAN Profile "via_corporate_wpa2") #ssid-profile "via_corporate_ssid"
```

For detailed configuration parameter information, see *"wlan client-wlan-profile" command* in the *Dell Networking W-Series ArubaOS 6.4.x Command Line Interface User Guide*.

Customize VIA logo, landing page and downloading installer

This step can only be performed using the WebUI. See Rebranding VIA and Downloading the Installer on page 694

Downloading VIA

This section of the document provides instructions and information on using VIA.

Pre-requisites

Ensure that the end-user system meets the following prerequisites:

- VIA can be installed only on systems running:
  - Microsoft Windows XP with SP2
  - Microsoft Windows Vista (32-bit and 64-bit)
  - Microsoft Windows 7 (32-bit and 64-bit)



VIA is supported only in the English versions of Microsoft Windows. International versions of Microsoft Windows is not supported.

- Requires the following Microsoft KB on the end-user systems:
  - On Microsoft Windows XP SP2—KB918997 (support.microsoft.com/kb/918997)

Install this to see the list of detected wireless networks in the VIA client (**Diagnostics** tab > **Detected Networks** page).

- On Microsoft Windows XP SP3—KB958071 (support.microsoft.com/kb/958071)

  Install this if you receive the "*1206 (ERROR_BAD_PROFILE)*" error code.

- Administrator rights on the computer.
- The computer must have a working wired or wireless network hardware.

Downloading VIA

In a typical scenario, end users will receive an email from their IT department with details to download VIA from a URL (controllers public IP address). See Table 130.

In this example, they can download VIA set up files from *https://115.52.100.10/via* after entering their corporate credentials.

Figure 99 *Login to Download VIA*

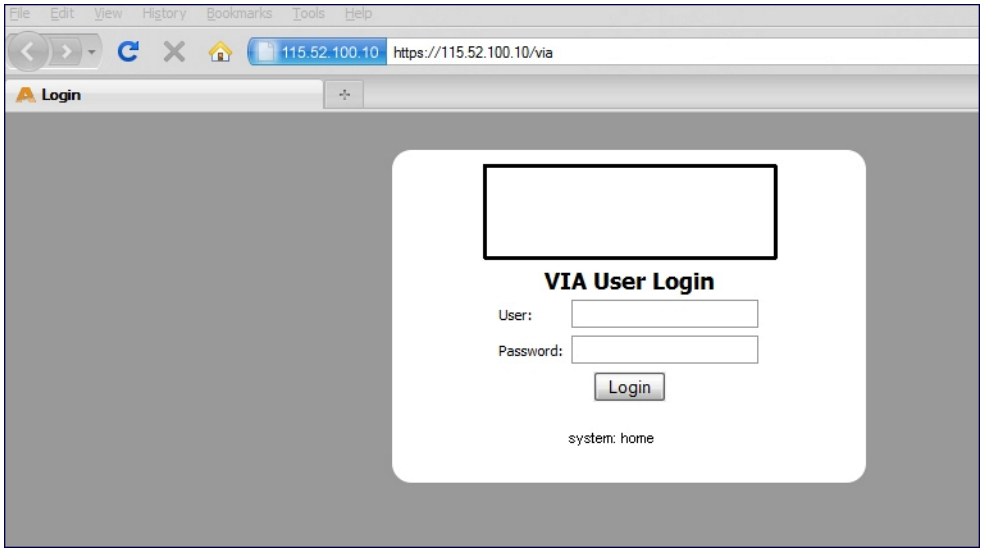
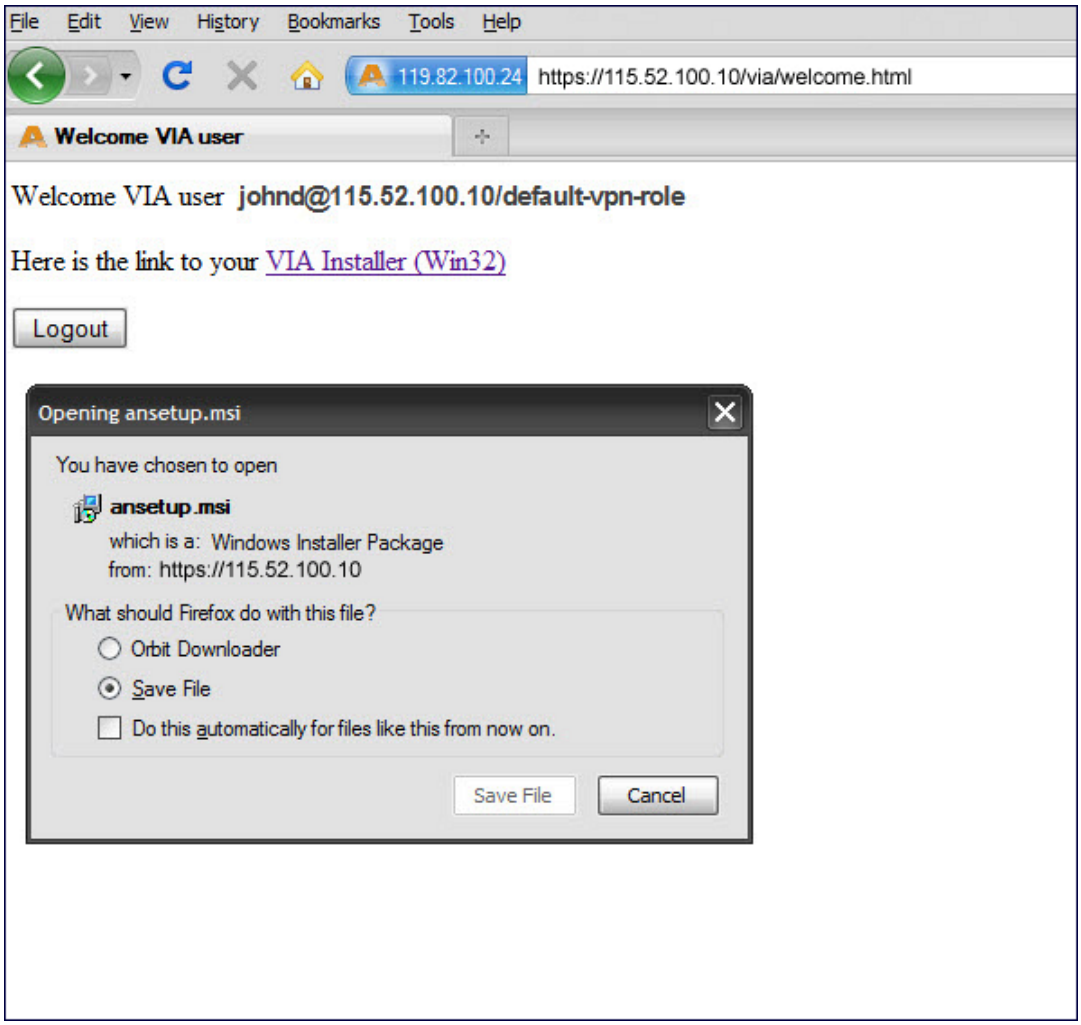


Figure 100 *Downloading VIA set up file after authentication*



Installing VIA

Double click the downloaded set up file (*ansetup.msi or ansetup64.msi*) to start the installation process. Ensure that you have met the pre-requisites before proceeding with the installation.

Using VIA

The VIA desktop application has three tabs:

- Connection Details
- Diagnostics
- Settings

Connection Details Tab

Provides all required details about your remote connection. After a successful connection, you can see the assigned IP from your remote server, the profile used for the connection and other network related information.

- Disconnect—Click this button to disconnect the current remote connection. You will have to manually connect for the next connection. VIA will not automatically start connection.

- View Connection Log—Click this button to view the sequence of events that took place during the last or current connection. The log also provide information about upgrade requirement, missing pre-requisites or other encountered errors.
- Change Profile—Click this button to select an alternate connection profile. This button is enabled only if your administrator has configured more than one connection profile. This button toggles to **Download Profile**, if you clear your profile from the Settings tab.

More Details

This section gives information about your local connection.

- Click**Network Details** to view local network connection information.
- Click**VIA Details** to view error or other connection messages.

Diagnostic Tab

Provides information and tools for troubleshooting your connectivity issues. Select a diagnostic tool from this tab for more information.

- Connection Logs—Sequence of events that happened during the recent connection.
- Send Logs—List of logs files collected by VIA. You can send this to your technical support when required. Click **Open Folder** to see the folder with the most recent logs and click the **Send** button to send log files archive using your default e-mail client.
- View system info & Advanced info—System and network configuration details of your system.
- Connectivity tests—Basic tests (ping and trace-route) to verify your network connection.
- Detected Networks—If your system has wireless network capability, this option will show all detected wireless networks.
- VIA info—Information about the current VIA installation.
- Compatibility info—Compatibility information about some applications detected in your system.

Settings Tab

This tab allows you to configure extra settings required to collect log, use a different connection profile and set up proxy server details.

- Log Settings—Allows you to set VIA log levels. By default, the log level is set to *Trace*. This setting captures extensive activity information about VIA.
- Connection Profile—Allows you to select and connect to a different connection profile. This is usually useful if you are in remote location and you need to connect to your corporate (secure) network. In such situation, you can select a profile that uses the nearest remote server to provide secure connection to your network. Alternate connection profiles are available only if it is configured by your IT administrator.
- Proxy Settings—Detects and displays Microsoft Internet Explorer proxy server details. It also allows you to enter the proxy authentication credentials to be used for HTTP/HTTPS connection to the controller.

Troubleshooting

To enable your support team to effectively resolve your VIA connection issues, it is mandatory that you send logs generated by VIA. To do this, click the **Send Logs**button from the **Connection Details** tab.

Chapter 32
Spectrum Analysis

Wireless networks operate in environments with electrical and radio frequency devices that can interfere with network communications. Microwave ovens, cordless phones, and even adjacent Wi-Fi networks are all potential sources of continuous or intermittent interference. The spectrum analysis software modules on APs that support this feature examine the radio frequency (RF) environment in which the Wi-Fi network is operating, identify interference and classify its sources. An analysis of the results quickly isolate issues with packet transmission, channel quality, and traffic congestion caused by contention with other devices operating in the same band or channel.

AP radios that gather spectrum data but do not service clients are called spectrum monitors, or SMs. Each SM scans and analyzes the spectrum band used by the SM's radio (2.4Ghz or 5Ghz). An AP radio in *hybrid AP* mode continues to serve clients as an access point while analyzing spectrum analysis data for the channel the radio uses to serve clients. You can record data for both types of spectrum analysis devices, save that data, and then play it back for later analysis.

Topics in this chapter include:

- Understanding Spectrum Analysis on page 700
- Creating Spectrum Monitors and Hybrid APs on page 705
- Connecting Spectrum Devices to the Spectrum Analysis Client on page 707
- Configuring the Spectrum Analysis Dashboards on page 710
- Customizing Spectrum Analysis Graphs on page 713
- Working with Non-Wi-Fi Interferers on page 739
- Understanding the Spectrum Analysis Session Log on page 741
- Viewing Spectrum Analysis Data on page 741
- Recording Spectrum Analysis Data on page 742
- Troubleshooting Spectrum Analysis on page 745

Understanding Spectrum Analysis

The table below lists the AP models that support the spectrum analysis feature. Single-radio mesh APs do not support the spectrum analysis feature; if an AP radio has a virtual AP carrying mesh backhaul traffic, no other virtual AP on that radio can be configured as a spectrum monitor. However, dual-radio mesh APs can have the client access radio configured as a Spectrum monitor or hybrid AP while the other radio supports mesh backhaul traffic.

Table 132: *Device Support for Spectrum Analysis*

Device	Configurable as a Spectrum Monitor?	Configurable as a Hybrid AP?
W-AP220 Series	Yes	Yes
W-AP114	Yes	Yes
W-AP115	Yes	Yes
W-AP104	Yes	Yes
W-AP105	Yes	Yes
W-AP92	Yes	Yes
W-AP93	Yes	Yes
W-AP93H	Yes	No
W-AP130 Series	Yes	Yes
W-AP175	Yes	No
W-IAP3WN Series	Yes	No
W-AP270 Series	Yes	Yes

The radios on groups of APs can be converted to dedicated spectrum monitors or hybrid APs via the AP group's dot11a and dot11g radio profiles. Individual APs can also be converted to spectrum monitors through the AP's spectrum override profile.



The spectrum analysis feature requires the RF Protect license. To convert an AP to a spectrum monitor or hybrid AP, you must have an AP license *and* an RFProtect license for each AP on that controller.

The Spectrum Analysis section of the **Monitoring** tab in the WebUI includes the **Spectrum Monitors**, **Session Log**, and **Spectrum Dashboards** windows.

- **Spectrum Monitors**: this window displays a list of active spectrum monitors and hybrid APs streaming data to your client, the radio band the device is monitoring, and the date and time the SM or hybrid AP was connected to your client. This window allows you to select the spectrum monitors or hybrid APs for which you want to view information, and release the connection between your client and any device you no longer want to view.
- **Session Log**: this tab displays activity for spectrum monitors and hybrid APs during the current browser session, including timestamps showing when the devices were connected to and disconnected from the client, and any changes to a hybrid APs monitored channel.
- **Spectrum Dashboards**: this window shows different user-customizable data charts for 2.4 Ghz and 5 GHz spectrum monitor or hybrid AP radios. Table 133 below gives a basic description of each of the spectrum analysis graphs that can appear on the spectrum dashboard.



For more detailed information on these graphs, refer to Customizing Spectrum Analysis Graphs on page 713.

Table 133: *Spectrum Analysis Graphs*

Graph Title	Description	Update Interval
Active Devices Table	A pie chart showing the percentages and total numbers of each device type for all active devices. This graph has no set update interval; the graph automatically updates when values change. For details, see Active Devices on page 714.	N/A
Active Devices Trend	A line chart showing the numbers of up to five different types of Wi-Fi and non-Wi-Fi devices seen on selected channels during a specified time interval. This chart can show devices on multiple channels for a spectrum monitor, or the single monitored channel for a hybrid AP. For details, see Active Devices Trend on page 718.	5 seconds
Channel Metrics	This stacked bar chart shows the current relative quality, availability or utilization of selected channels in the 2.4 GHz or 5 GHz radio bands. This chart can show multiple channels for a spectrum monitor, or the single monitored channel for a hybrid AP. For details, see Channel Metrics on page 719.	5 seconds
Channel Metrics Trend	A line chart showing the relative quality or availability of selected channels in the 2.4 GHz or 5 GHz radio bands over a specified time interval. Spectrum monitors can show channel data for multiple channels, while a hybrid AP shows information only for its one monitored channel. For details, see Channel Metrics Trend on page 721.	5 seconds
Channel Summary Table	The Channel Summary table displays the number of devices found on each channel in the spectrum monitor's radio band, the percentage of channel utilization, and AP power and interference levels. Spectrum monitors can show data for multiple channels, while a hybrid AP shows a channel summary only for its one monitored channel. For details, see Channel Summary Table on page 723.	5 seconds
Channel Utilization Trend	A line chart that shows the channel utilization for one or more radio channels, as measured over a defined time interval. Spectrum monitors can show data for multiple channels, while a hybrid AP shows utilization levels for its one monitored channel only. For details, see Channel Utilization Trend on page 726.	5 seconds
Device Duty Cycle	A stacked bar chart showing the percent of each channel in the spectrum monitor radio's frequency band used by a Wi-Fi AP or any other device type detected by the spectrum monitor. The Device Duty Cycle chart for a hybrid AP only shows data for the one channel monitored by the hybrid AP.This chart is not available for W-AP68 access points. For details, see Device Duty Cycle on page 724.	5 seconds
Devices vs Channel	A stacked bar chart showing the total numbers of each device type detected on each channel in the spectrum monitor radio's frequency band. The Devices vs Channel chart for a hybrid AP only shows data for the one channel monitored by the hybrid AP. For details, see Devices vs Channel on page 727.	5 seconds

Graph Title	Description	Update Interval
FFT Duty Cycle	Fast Fourier Transform, or **FFT**, is an algorithm for computing the frequency spectrum of a time-varying input signal. This line chart shows the FFT duty cycle, which represents the percent of time a signal is broadcast on the specified channel or frequency. Spectrum monitors can show data for multiple channels, while a hybrid AP shows information only for its one monitored channel. This chart is not available for W-AP68 access points. For details, see FFT Duty Cycle on page 729.	1 second
Interference Power	This chart shows information about Wi-Fi interference, including the Wi-Fi noise floor, and the amount of adjacent channel interference from cordless phones, bluetooth devices and microwaves. Spectrum monitors can show interference power data for multiple channels, while a hybrid AP shows information only for its one monitored channel. For details, see Interference Power on page 731.	5 seconds
Quality Spectrogram	This plot shows quality statistics for selected range of channels or frequencies as determined by the current noise floor, non-Wi-Fi (interferer) utilization and duty-cycles and certain types of retries. This chart can also be configured to show channel availability, the percentage of each channel that is unused and available for additional traffic. Spectrum monitors can show data for multiple channels, while a hybrid AP shows information only for its one monitored channel. For details, see Quality Spectrogram on page 733.	5 seconds
Real-Time FFT	Fast Fourier Transform, or **FFT**, is an algorithm for computing the frequency spectrum of a time-varying input signal. This line chart shows the power level of a signal on the channels or frequencies monitored by a spectrum monitor radio. Spectrum monitors can show data for multiple channels, while a hybrid AP shows information only for its one monitored channel. This chart is not available for W-AP68 access points. For details, see Real-Time FFT on page 734.	1 second
Swept Spectrogram	This plot displays FFT power levels For details, see or the FFT duty cycle for a selected channel or frequency, as measured during each time tick. Spectrum monitors can show data for multiple channels, while a hybrid AP shows information only for its one monitored channel. This chart is not available for W-AP68 access points. For details, see Swept Spectrogram on page 736.	1 second

Spectrum Analysis Clients

The maximum number of spectrum monitor radios and hybrid AP radios on a controller is limited only by the number of APs on that controller. If desired, you can configure every radio on an AP that supports the Spectrum Analysis feature as a spectrum device. A dual-radio AP can operate as two spectrum devices, because each radio can be individually configured as a spectrum monitor (SM) or hybrid AP.

A spectrum analysis client can simultaneously access data from up to four individual spectrum device radios. Each spectrum device radio, however, can only be connected to a single client WebUI.

When you select a specific spectrum monitor or hybrid AP radio to stream data to your client, the controller first verifies the device is not subscribed to some other client. Once the SM or hybrid AP radio has been verified as available, the SM or hybrid AP establishes a connection to the client and begins sending spectrum analysis

data either every second or every five seconds, depending on the type of data being requested. Each client may select up to twelve different spectrum analysis charts and graphs to appear in the spectrum dashboard.

A controller can support up to 22 active WebUI connections. If spectrum analysis clients are simultaneously viewing data for than 22 WebUI connections, any additional WebUI requests are refused until some clients close their WebUI browser sessions.

When you finish reviewing data from an SM or hybrid AP, you should disconnect the device from your spectrum client. Do not forget this important step—no other user can access data from that spectrum monitor or hybrid AP until you release your subscription. Note, however, that when you disconnect a spectrum monitor from your client, *the AP continues to operate as a spectrum monitor* until you return it to AP mode by removing the local spectrum override, or by changing the mode parameter in the AP's 802.11a or 802.11g radio profile from spectrum-mode back to AP-mode.



A spectrum monitor or hybrid AP automatically disconnects from a client when you close the browser window you used to connect the spectrum monitor to your client. However, if you use Internet Explorer and have multiple instances of an Internet Explorer browser open, the data-streaming connection to the spectrum monitor or hybrid AP is not released until 60 seconds after you close the spectrum client browser window. During this 60-second period, the spectrum monitor is still connected to the client.

When a spectrum monitor or hybrid AP is not subscribed to any client, it still performs all classification tasks and collect all necessary channel lists and device information. You can view classification, device, and channel information for any active spectrum monitor or hybrid AP via the controller's command-line interface, regardless of whether or not that device is sending real-time spectrum data to another client WebUI.

Individual spectrum analysis graphs and charts are explained in detail in Customizing Spectrum Analysis Graphs on page 713.

Hybrid AP Channel Changes

By default, a hybrid AP only monitors the channel specified in its 802.11a or 802.11g radio profile for spectrum interference. If you want to change the channel monitored by a hybrid AP, you must edit the channel setting in those profiles. However, there are other ArubaOS features that may automatically change the channels on hybrid APs. APs using Dynamic Frequency Selection (DFS) perform off-channel scanning to detect the presence of satellite and radar transmissions, and switch to a different channel if it detects that satellite or radar transmissions are present. APs using the Adaptive Radio Response (ARM) feature constantly monitor the network and automatically select the best channel and transmission power settings for that AP. If you manually change a channel monitored by a hybrid AP, best practices are to temporarily disable the ARM feature, as ARM may automatically return the channel to its previous setting.

If a hybrid AP is using ARM or DFS, that hybrid AP may automatically move to a different channel in response to changes in the network environment. If a hybrid AP changes channels while it is connected to a spectrum analysis client, the hybrid AP updates the graphs in the spectrum dashboard to start displaying spectrum data for the new channel, and sends a log message to the session log. For details on changing the channel monitored by a hybrid AP, see 802.11a and 802.11g RF Management Profiles on page 509.

Hybrid APs Using Mode-Aware ARM

If a radio is configured as a hybrid AP and that AP is enabled with mode-aware ARM, the hybrid AP can change to an Air Monitor (or *AM*) if too many APs are detected in the area. If the ARM feature changes a hybrid AP to an Air Monitor, that AM does not provide spectrum data after the mode change. The AM unsubscribes from any connected spectrum analysis client, and sends a log message warning about the change. If mode-aware ARM changes the AM back to an AP, the hybrid AP does not automatically resubscribe back to the spectrum analysis client. The hybrid AP must manually resubscribed before it can appear in the client's **spectrum monitors** page.

Creating Spectrum Monitors and Hybrid APs

Each controller can support up to 22 active WebUI connections to spectrum monitor or hybrid AP radios. If you plan on using spectrum monitors or hybrid APs as a permanent overlay to constantly monitor your network, you should create a separate AP group for these devices. If you plan on temporarily converting campus APs to spectrum monitors, best practices are to use the spectrum local override profile to convert an AP to a spectrum monitor.

This section describes the following tasks for converting regular APs into hybrid APs or spectrum monitors.

- Converting APs to Hybrid APs on page 705
- Converting an Individual AP to a Spectrum Monitor on page 706
- Converting a Group of APs to Spectrum Monitors on page 706

Converting APs to Hybrid APs

You can convert a group of regular APs into a hybrid APs by selecting the **spectrum monitoring** option in the AP group's 802.11a and 802.11g radio profiles. Once you have enab led the spectrum monitoring option, all APs in the group that support the spectrum monitoring feature start to function as hybrid APs. If any AP in the group does *not* support the spectrum monitoring feature, that AP continues to function as a standard AP, rather than a hybrid AP.



The spectrum monitoring option in the 802.11a and 802.11g radio profiles only affects APs in ap-mode. Devices in am-mode (Air Monitors) or sm-mode (Spectrum Monitors) are not affected by enabling this option.

If you want to convert a individual AP (and not an entire AP group) to a hybrid AP, you must create a new 802.11a or 802.11g radio profile, enable the **spectrum monitoring** option, then reassign that AP to the new profile. For additional information see Creating and Editing Mesh High-Throughput SSID Profiles on page 556 for details on how to create a new 802.11a/g radio profile, then assign an individual AP to that profile.



If the spectrum local-override profile on the controller that terminates the AP contains an entry for a hybrid AP radio, that entry overrides the mode selection in the 802.11a or 802.11g radio profile, and the AP operates as a spectrum monitor, *not* as a hybrid AP. You must remove any spectrum local override for an AP to allow the device to operate as a hybrid AP. For further details on editing a spectrum local override, see Converting an Individual AP to a Spectrum Monitor on page 706.

In the WebUI

Follow the procedure below to convert a group of APs to hybrid mode via the WebUI.

1. Navigate to the **Configuration** > **Wireless** > **AP Configuration** window. Select the **AP Group** tab.
2. Click **Edit** by the name of the AP group you want to convert to hybrid APs.
3. Under the **Profiles** list, expand the **RF Management** menu.
4. To enable a spectrum monitor on the 802.11a radio band, select the **802.11a radio profile** menu.

   -or-

   To enable a spectrum monitor on the 802.11g radio band, select the **802.11g radio profile** menu.
5. The **Profile Details** pane appears. Select the **Spectrum Monitor** checkbox.
6. Click **Apply**.

In the CLI

To convert a group of APs via the command-line interface, access the CLI in config mode and issue the following commands, where `<profile>` is the name of the 802.11a or 802.11g radio profile used by the group of APs you want to convert to hybrid APs:

```
rf dot11a-radio-profile <profile> spectrum-monitoring
rf dot11g-radio-profile <profile> spectrum-monitoring
```

Converting an Individual AP to a Spectrum Monitor

There are two ways to change a radio on an individual AP or AM into a spectrum monitor. You can assign that AP to a different 802.11a and 802.11g radio profile that is already set to spectrum mode, or you can temporarily change the AP into a spectrum monitor using a local spectrum override profile. When you use a local spectrum override profile to override an AP's mode setting, that AP begins to operate as a spectrum monitor, but remains associated with its previous 802.11a and 802.11g radio profiles. If you change any parameter (other than the overridden **mode** parameter) in the spectrum monitor's 802.11a or 802.11 radio profiles, the spectrum monitor immediately updates with the change. When you remove the local spectrum override, the spectrum monitor reverts back to its previous mode, and remains assigned to the same 802.11a and 802.11 radio profiles as before.

The spectrum local override profile overrides the **mode** parameter in the 802.11a or 802.11g radio profile, changing it from ap-mode or am-mode to spectrum-mode, while allowing the spectrum monitor to continue to inherit all other settings from its 802.11a/802.11g radio profiles. When the spectrum local override is removed, the AP automatically reverts to its previous mode as defined it its 802.11a or 802.11g radio profile settings. If you use the local override profile to change an AP radio to a spectrum monitor, you must do so by accessing the WebUI or CLI of the controller that terminates the AP. This is usually a local controller, not a master controller.

In the WebUI

To convert an individual AP using the local spectrum override profile in the WebUI:

1. Select **Configuration** > **All Profiles**. The **All Profile Management** window opens.
2. Select **AP** to expand the **AP** profiles section.
3. Select **Spectrum Local Override Profile**. The **Profile Details** pane displays the current **Override Entry** settings.
4. In the **AP name** entry blank, enter the name of an AP whose radio you want to configure as a spectrum monitor. Note that AP names are case-sensitive. Any extra spaces before or after the AP name prevents the AP from being correctly added to the override list.
5. If your AP has multiple radios or a single dual-band radio, click the **band** drop-down list and select the spectrum band you want that radio to monitor: **2-ghz** or **5-ghz**. Click **Add** to add that radio to the **Override Entry** list.
6. (Optional) Repeat steps 4-6 to convert other AP radios to spectrum monitors, as desired. To remove a spectrum monitor from the override entry list, select that radio name in the override entry list, then click **Delete**.
7. Click **Apply**.

In the CLI

To convert an individual AP spectrum monitor using the spectrum local override profile in the command-line interface, access the CLI in config mode and issue the following command:

```
ap spectrum local-override override ap-name <ap-name> spectrum-band 2ghz|5ghz
```

Converting a Group of APs to Spectrum Monitors

When you convert a group of APs to spectrum monitors using their 802.11a/802.11g radio profiles, all AP radios associated with that profile stop serving clients and act as spectrum monitors only. Therefore, before you convert an entire group of APs to spectrum monitors, be sure that none of the APs are currently serving clients, as that may temporarily interrupt service to those clients.



If you use an 802.11a or 802.11g radio profile to create a group of spectrum monitors, all APs in any AP group referencing that radio profile are set to spectrum mode. Therefore, best practices are to create a new 802.11a or 802.11g radio profile just for spectrum monitors, using the following CLI commands:

```
ap-name <ap name> dot11a-radio-profile <profile-name>ap-name <ap name> dot11g-radio-profile <profile-name>
```

If you want to set an existing 802.11a or 802.11g radio profile to spectrum mode, verify that no other AP group references that radio profile, using the following CLI commands:

```
show references rf dot11a-radio-profile <profile-name>show references rf dot11g-radio-profile <profile-name>
```

In the WebUI

Follow the procedure below to convert a group of APs to Spectrum mode via the WebUI.

1. Navigate to the **Configuration** > **Wireless** > **AP Configuration** window. Select the **AP Group** tab.
2. Click **Edit** by the name of the AP group you want to convert to spectrum monitors.
3. Under the **Profiles** list, expand the **RF Management** menu.
4. To enable a spectrum monitor on the 802.11a radio band, select the **802.11a radio profile** menu.

   -or-

   To enable a spectrum monitor on the 802.11g radio band, select the **802.11g radio profile** menu.
5. The **Profile Details** pane appears. Click the **Mode** drop-down list, and select **spectrum-mode**.
6. Click **Apply**.

In the CLI

To convert a group of APs via the command-line interface, access the CLI in config mode and issue the following commands, where **<profile>** is the 80211a or 80211g radio profile used by the AP group.

```
rf dot11a-radio-profile <profile> mode spectrum-mode
rf dot11g-radio-profile <profile> mode spectrum-mode
```

Connecting Spectrum Devices to the Spectrum Analysis Client

A spectrum analysis client is any laptop or desktop computer that can access the controller WebUI and receive streaming data from individual spectrum monitors or hybrid APs. Once you have configured one or more APs to operate as a spectrum monitor or hybrid AP, use the **Spectrum Monitors** window to identify the spectrum devices you want to actively connect to the spectrum analysis client. To connect one or more spectrum devices to your client:

1. Navigate to **Monitoring** > **Spectrum Analysis**.
2. Click the **Spectrum Monitors** tab.
3. Click **Add**. A table appears, displaying a list of spectrum analysis devices, sorted by name. Single-radio spectrum devices have a single entry in this table, and dual-radio spectrum devices have two entries: one for each radio. This table displays the following data for each radio.

Table 134: *Spectrum Device Selection Information*

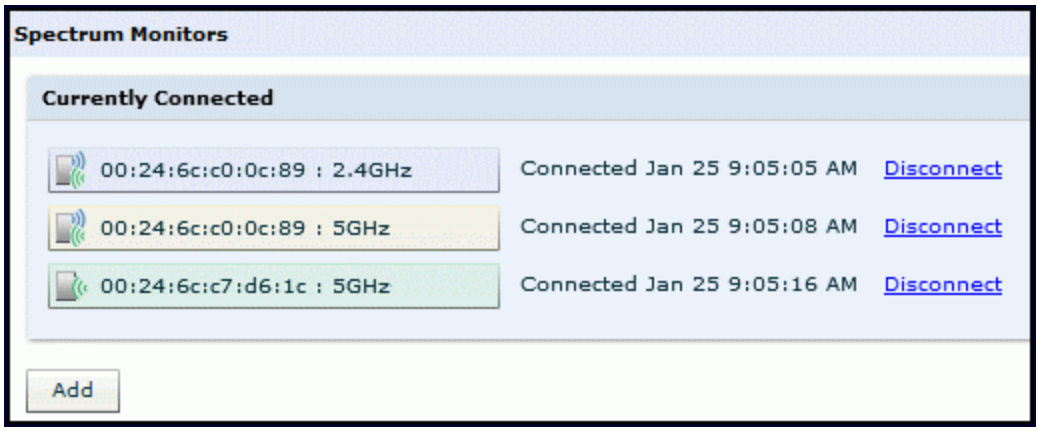
Table Column	Description
AP	Name of the AP whose radio you want to convert to a spectrum monitor. AP names are case sensitive. This column includes the following icons:  Radio is operating as a spectrum monitor.  Radio is operating as a hybrid AP with spectrum enabled.
Band	The frequency band currently used by the radio. This value can be either **2.4 GHz** or **5 GHz**.
Model	AP model type.
AP Group	Name of the AP group to which the spectrum monitor is currently associated.
Mode	This column indicates the type of spectrum analysis device: • **Spectrum Monitor**: AP is in spectrum monitor mode. • **Access Point**: AP is configured as an access point but with spectrum monitoring enabled (Hybrid AP).
Availability for Connection	Indicates if the AP is available to send spectrum analysis data to the client. Possible options ares: • **Available, 2.4GHz**: the radio is available to send spectrum analysis data on the 2.4GHz frequency band. • **Available, 5GHz**: the radio is available to send spectrum analysis data on the 5GHz frequency band. • **Available, Dual Band**: the radio is available and is capable of sending spectrum analysis data on either the 2.4 GHz or the 5 GHz frequency bands. • **Available, current channel - <channel>**: the AP radio is in hybrid mode and can display spectrum analysis data for the single specified channel only. • **Not available:** an AP may not be available because it is currently sending spectrum analysis data to another client.

4. Click the table entry for a spectrum monitor radio, then click **Connect**.
5. Repeat steps 3-4 to connect additional devices, if desired.

View Connected Spectrum Analysis Devices

Once you have connected one or more spectrum monitors or hybrid APs to your Spectrum Analysis client, the **Monitoring** > **Spectrum Analysis** > **Spectrum Monitors** window displays a table of currently connected spectrum devices. This table includes the name of each spectrum monitor or hybrid AP and its current radio band (2GHz or 5GHz):

Figure 101 *Viewing a list of Connected Spectrum Monitors*



To view a list of connected spectrum devices via the command-line interface, issue the **show ap spectrum monitors** command:

```
(host)# show ap spectrum monitors
List of Sensors
---------------
AP name  Group    AP Type  Phy  Band       Channel  Mode             Client IP
Subscribe Time
-------  -----    -------  ---  ----       -------  ----             ---------
--------------
AP12     default  105      G    2GHz       -        Spectrum Monitor 10.4.165.227
2011-04-25 02:53:52 AM
AP16     default  135      A    5GHz       149      Access Point     10.4.165.227
2011-04-25 02:53:55 AM
AP44     default  105      G    2GHz       -        Spectrum Monitor 10.4.165.227
2011-04-25 02:54:03 AM
```

Disconnecting a Spectrum Device

A spectrum monitor or hybrid AP can send spectrum analysis data to only one client at a time. When you are done viewing data for a spectrum device, you should release your client's subscription to that spectrum device and allow other clients to view data from that device. A spectrum monitor or hybrid AP automatically disconnects from your client when you close the browser window used to connect the spectrum device your client.

To manually disconnect a spectrum monitor or hybrid AP:

1. Click the **Spectrum Monitors** tab.
2. Each table entry in the **Currently Connected** table includes a **Disconnect** link to release the client's connection to that spectrum monitor. Identify the table entry for the spectrum monitor you want to release then click **Disconnect**.
3. A popup window asks you to confirm that you want to disconnect the spectrum monitor from the spectrum analysis client. Click **OK**. The spectrum monitor d>isconnects from the client and the device's entry is removed from the **Currently Connected** table.

When you disconnect a spectrum device from your client, the AP continues to operate as a spectrum monitor or hybrid AP until you return the device to AP mode by removing the local spectrum override, or by changing the mode parameter in the AP's 802.11a or 802.11g radio profile from spectrum-mode to AP-mode.



If you are use Internet Explorer with multiple instances of the Internet Explorer browser open, and you close the spectrum browser window without manually disconnecting the spectrum device, the controller does not release the data streaming connection to aspectrum monitor until 60 seconds after you close the spectrum client browser window. During this 60-second period, the spectrum monitor is still connected to the client.

Configuring the Spectrum Analysis Dashboards

Once you have connected spectrum monitors to your spectrum analysis client, you can begin to monitor spectrum data in the spectrum analysis dashboards. There are three predefined sets of dashboard views, **View 1**, **View 2** and **View 3**. View 1 displays the Real-Time FFT, FFT Duty-Cycle and Swept Spectrogram graphs by default, and Views 2 and 3 display the Swept Spectrogram and Quality Spectrogram charts, and the Channel Summary and Active Devices tables.

Each chart in the dashboard can be replaced with other chart types, or reconfigured to show data for a different spectrum monitor. Once you have configured a dashboard view with different settings, you can rename that dashboard view to better reflect its new content.

The following sections explain how to customize your Spectrum Analysis dashboard to best suit the needs of your individual network:

- Selecting a Spectrum Monitor on page 710
- Changing Graphs within a Spectrum View on page 711
- Renaming a Spectrum Analysis Dashboard View on page 711
- Saving a Dashboard View on page 712
- Resizing an Individual Graph on page 713

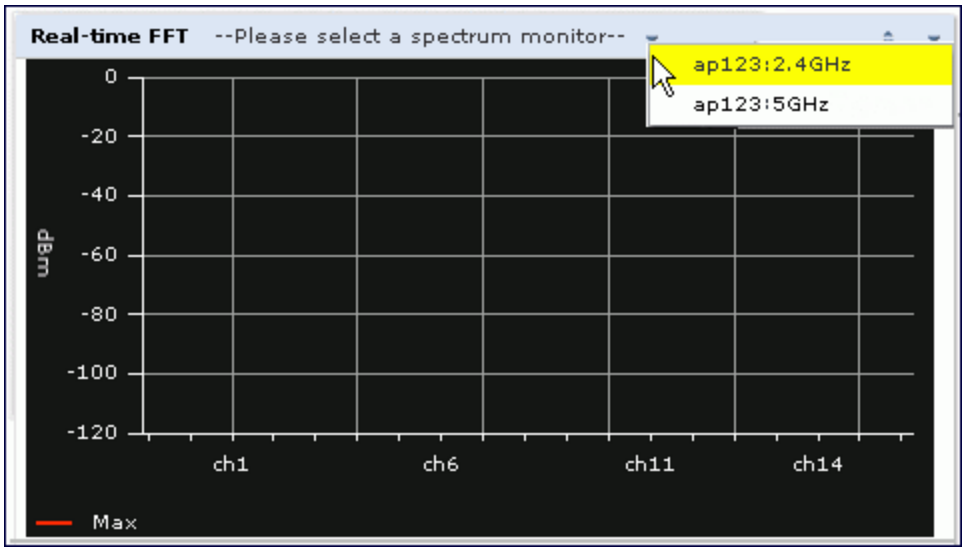
Selecting a Spectrum Monitor

When you first log into the Spectrum Analysis dashboard, it displays blank charts. You must identify the spectrum monitor whose information you want to view before the graphs display any data.

To identify the spectrum monitor radio whose data you want to appear in the Spectrum Analysis dashboard:

1. Access the **Monitoring** > **Spectrum Analysis** window in the WebUI.
2. Click the **Spectrum Dashboards** tab.
3. In the graph title bar, click the down arrow by the **Please select a spectrum monitor** heading, as shown in Figure 102. A drop-down list appears with the name of all spectrum monitor and hybrid AP radios currently connected to the client.

Figure 102 *Selecting a Spectrum Monitor*



4. Select a spectrum monitor from the list. The spectrum monitor or hybrid AP name appears in the chart titlebar and the chart starts displaying data for that spectrum monitor.

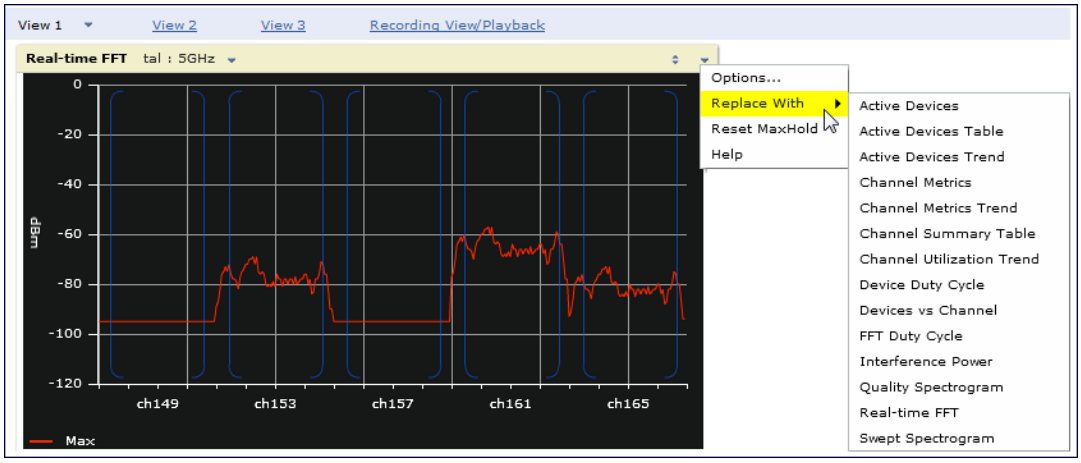
After you have selected the initial spectrum monitor or hybrid AP for a graph, you can display data for a different spectrum device at any time by clicking the down arrow by the device name in the chart titlebar and selecting a different connected spectrum monitor or hybrid AP.

Changing Graphs within a Spectrum View

To replace an existing graph with any other type of graph or chart:

1. Access the **Monitoring > Spectrum Analysis** window in the WebUI.
2. Click the **Spectrum Dashboards** tab.
3. From **Spectrum Dashboards** window, click one of the view names at the top of the window to select the dashboard layout with the graph you want to change.
4. Click the down arrow at the far right end of the graph title bar to display a drop-down menu of chart options.
5. Click **Replace With** to display a list of available graphs.
6. Click the name of the new graph you want to display.

Figure 103 *Replacing a Graph in the Spectrum Analysis Dashboard*



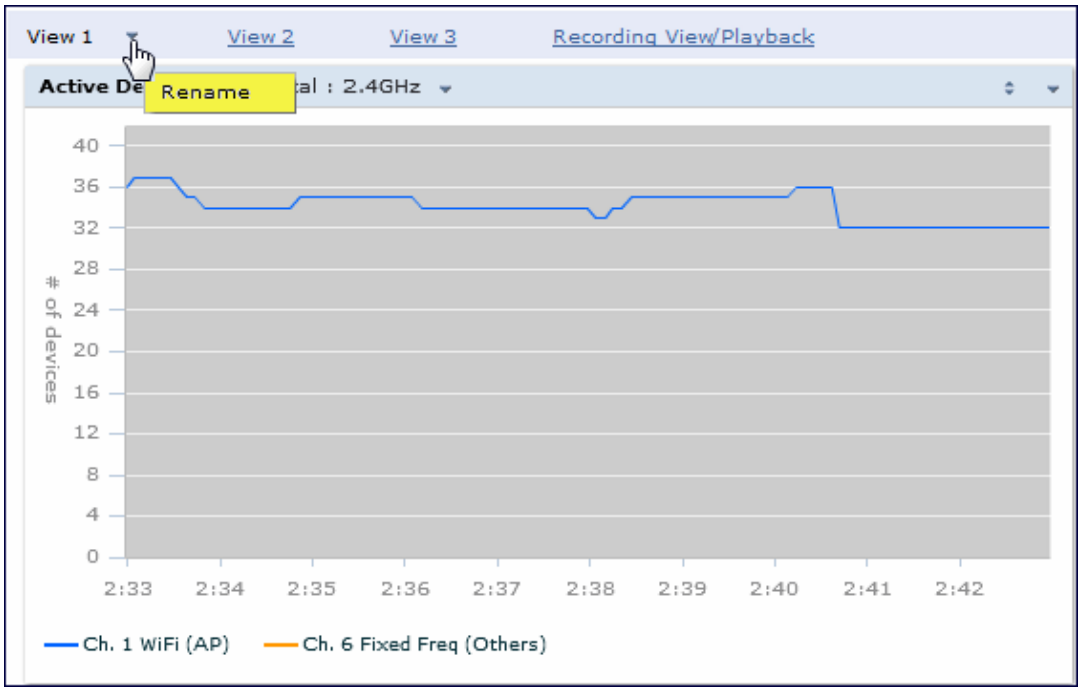
Renaming a Spectrum Analysis Dashboard View

You can rename any of the three spectrum analysis dashboard views at any time. Note, however, that simply renaming a view does not save its settings. (For details on saving a spectrum dashboard view, refer to Saving a Dashboard View on page 712.)

To rename a Spectrum Analysis Dashboard view:

1. From the **Monitoring > Spectrum Analysis > Spectrum Dashboards** window, click the down arrow to the right of the dashboard view you want to rename.
2. Select **Rename**.

Figure 104 *Renaming a Spectrum Dashboard View*



3. The **Dashboard Name** popup window appears. Enter a new name for the dashboard view, then click **OK**.

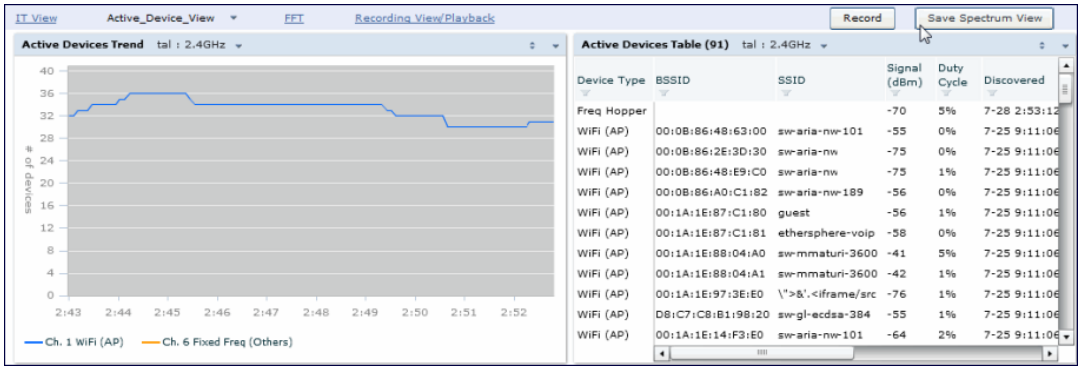
Saving a Dashboard View

You can select different graphs to display in a dashboard view, but these changes are not saved unless you save that view. Dashboard views, (like the spectrum analysis profile and spectrum local-override profile) are all local configurations that must be configured on each controller. None of these settings are synchronized between controllers.

To save a dashboard view:

1. After selecting the graphs you want to appear in the view, click **Save Spectrum View** at the top of the window.

Figure 105 *Save a Spectrum Analysis Dashboard Layout*



2. The **Spectrum View Saved** confirmation window appears when the spectrum view has been saved. The selected graphs now appear by default whenever you log in to view the spectrum dashboard.

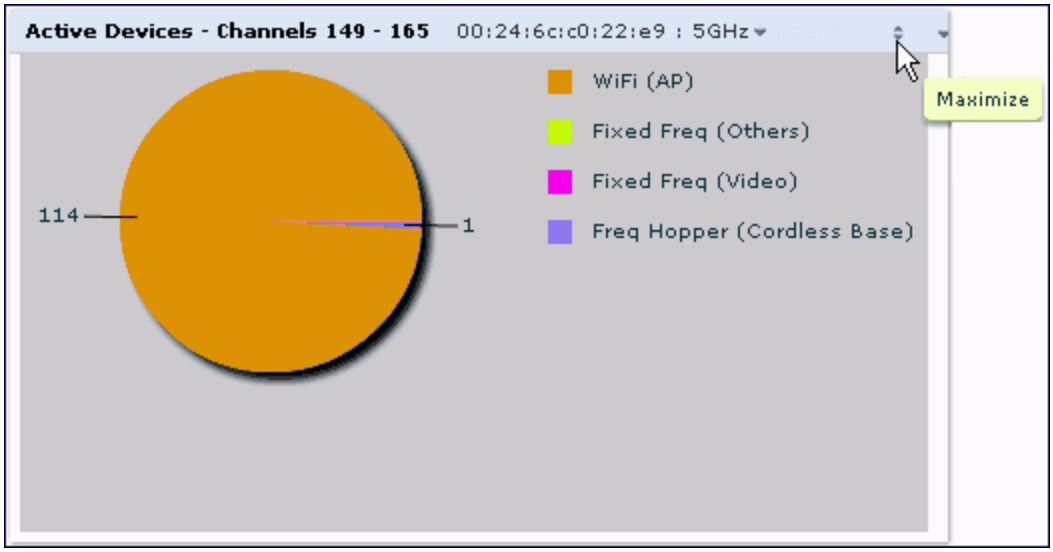
NOTE

If you change graphs in a spectrum view but do not save your settings, you are prompted to save or cancel your changes when you close the spectrum dashboard browser window

Resizing an Individual Graph

The left side of the title bar for each graph includes a resizing button on that allows you to expand a graph for easier viewing. Click this button as shown in Figure 106 to expand the selected graph to the size of the full window and display the **Options** pane, which allows you to change the current display options for that graph. (Configuration options are described in Spectrum Analysis Graph Configuration Options on page 714). To close the options pane if you have not made any changes to the graph, click **Close** at the bottom of the **Options** pane *or* click the resize button again to return the graph to its original size. To save any changes to the graph, click **OK** to save your settings and close the **Options** pane.

Figure 106 *Resizing a Spectrum Analysis Graph*

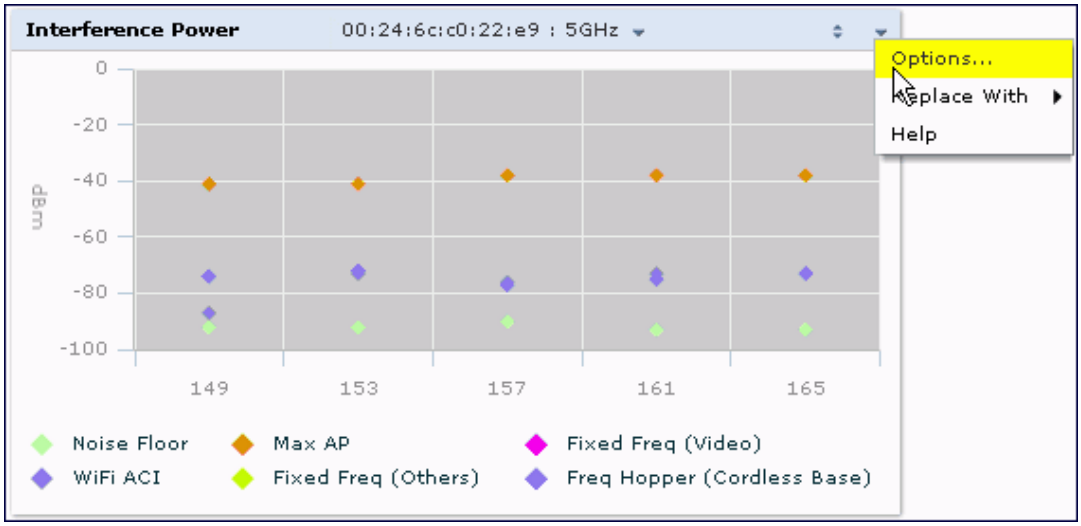


Customizing Spectrum Analysis Graphs

Each Spectrum Analysis graph can be customized to display or hide selected data types. To view the available options for a graph type:

1. From the **Monitoring > Spectrum Analysis > Spectrum Dashboards** window, click the down arrow at the end of the title bar for the graph you want to configure.
2. Select **Options**. The **Options** window appears to the right of the graph.

Figure 107 *Viewing Spectrum Analysis Graph Options*



3. From the **Options** window, configure graph settings described in Spectrum Analysis Graph Configuration Options on page 714.
4. When you are done, click **OK** at the bottom of the **Options** window to hide the options window.
5. (Optional) Click **Save Spectrum View** at the top of the window to save your new settings.

Spectrum Analysis Graph Configuration Options

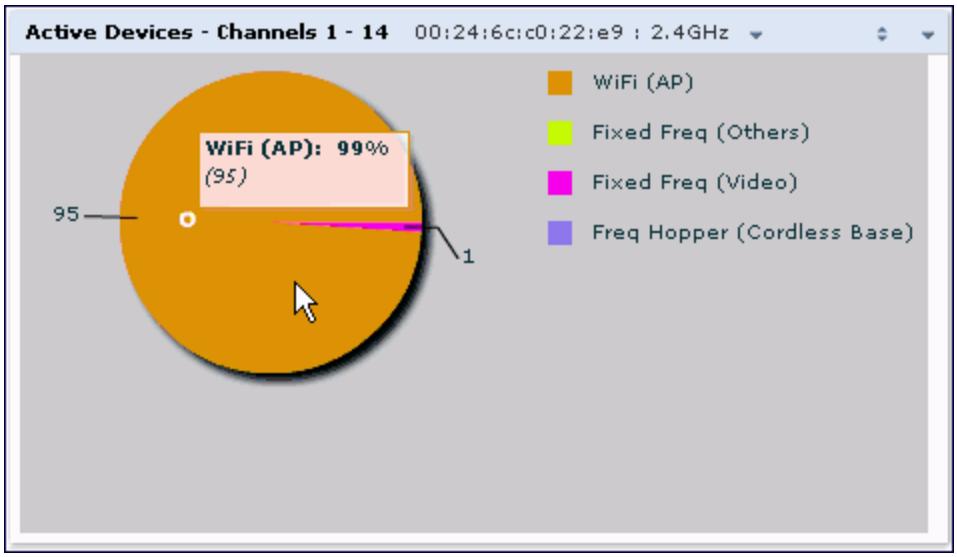
The following sections describe the customizable parameters and the default settings for each spectrum analysis graph.

Active Devices

This graph appears as a pie chart showing the percentages and total numbers of each device type for all active devices seen by the spectrum monitor or hybrid AP radio. This chart is useful for determining which types of devices are sending signals on the specified radio band or channel. The Active Devices graphs for spectrum monitors can be configured to show data for several different device types on a single radio channel or range of channels. Active Devices graphs for hybrid APs can show data for the single monitored channel only.

When you hover your mouse over any section of the pie chart, a tooltip displays the percentage and number of active devices classified into that device type. The example in Figure 108 shows that 99% of the active devices a spectrum monitor radio sees in the 2.4 GHz band are Wi-Fi APs.

Figure 108 *Active Devices Graph*



Click the down arrow in the upper right corner of this chart then click the **Options** menu to access the configuration settings for the Active Devices graph. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

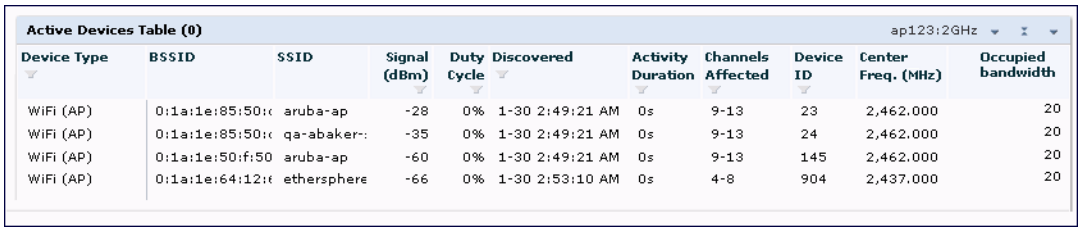
Table 135: *Active Devices Graph Options*

Parameter	Description
Band	Radio band displayed in this graph (2.4 GHz or 5 GHz)
Channel numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.
Channel Range	For graphs created by spectrum monitors, specify a channel range to determine which channels appears in this graph. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the graph. This graph displays all channels within the spectrum monitor's radio band by default. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.
Show	Click the checkbox by any of these device categories to include that device type in the graph. • WiFi (AP) • Microwave *(This option is only available for 2.4 GHz radios)* • Bluetooth *(This option is only available for 2.4 GHz radios)* • Fixed Freq (Others) • Fixed Freq (Cordless Phones) • Fixed Freq (Video) • Fixed Freq (Audio) • Freq Hopper (Others) • Freq Hopper (Cordless Network) • Freq Hopper (Cordless Base) • Freq Hopper Xbox *(This option is only available for 2.4 GHz radios)* • Microwave (Inverter) *(This option is only available for 2.4 GHz radios)* • Generic Interferer **NOTE:** For more information on non-Wi-Fi device types detected by a spectrum monitor, see Working with Non-Wi-Fi Interferers on page 739.

Active Devices Table

This table lets you view, sort, and search for data about the devices that are sending signals on the specified radio band or channel. The Active Devices table for a spectrum monitor displays data for all channels on the selected band. The Active Devices table for a hybrid monitor displays data for the single monitored channel only. Click any of the column headings to sort the information in the table by that column criteria. Make a column wider or narrower by clicking the border of a column heading and dragging the border to a new position.

Figure 109 *Active Devices Table*



You can save the data in the Active Devices table for later analysis by exporting it as data file in .csv format, which can be viewed by spreadsheet and database management applications like Microsoft Excel. To export this table, click the down arrow in the upper right corner of this chart and select **Export**. A window opens and lets you browse to the location to which you want to save the file. Once you have identified the location where you want to save the file, click **Save**.

You can also filter table entries by signal strength, duty cycle, discovery time, activity duration, channels affected and device ID number by clicking the icon below any column heading and specifying the values or value ranges that should appear in the table. Table 136 describes each of the columns in the table and the filters that can be applied to the table output.

Table 136: *Active Devices Table Options*

Parameter	Description
Device Type	This column shows the type of active device detected by the spectrum monitor or hybrid AP. This column may display any of the following values: • WiFi (AP) • Microwave *(This option is only available for 2.4 GHz radios)* • Bluetooth *(This option is only available for 2.4 GHz radios)* • Fixed Freq (Others) • Fixed Freq (Cordless Phones) • Fixed Freq (Video) • Fixed Freq (Audio) • Freq Hopper (Others) • Freq Hopper (Cordless Network) • Freq Hopper (Cordless Base) • Freq Hopper Xbox *(This option is only available for 2.4 GHz radios)* • Microwave (Inverter) *(This option is only available for 2.4 GHz radios)* • Generic Interferer **NOTE:** For more information on non-Wi-Fi device types detected by a spectrum monitor, see Working with Non-Wi-Fi Interferers on page 739.
BSSID	The Basic Service Set Identifier of the device. An AP's BSSID is usually its MAC address.
SSID	The service set identifier of the device's 802.11 wireless LAN.
Signal (dBm)	The current transmission power for this device, in dBm. To filter the output of this table to show only specific device types, click the icon in the column heading then select one of the following options: • Select **Any** to display entries for all signal strength levels. • To display entries within a specific range of power strength levels, enter the minimum signal strength level in the **Min** field and enter the maximum signal strength level in the **Max** field. Click **OK** to save your settings and return to the **Active Devices** table.
Duty Cycle	The percentage of time that the device is actively sending a signal on the radio band or channel. To filter the output of this table to show only specific duty cycle values or a range of values types, click the icon in the column heading then select one of the following options: • Select **Any** to display all entries, regardless of duty cycle value. • To display entries within a specific range of duty cycles, enter the minimum duty cycle percentage in the **Min** field and enter the maximum duty cycle percentage in the **Max** field. Click **OK** to save your settings and return to the **Active Devices** table.
Discovered	The time at which the device was first discovered by the spectrum monitor or hybrid AP.

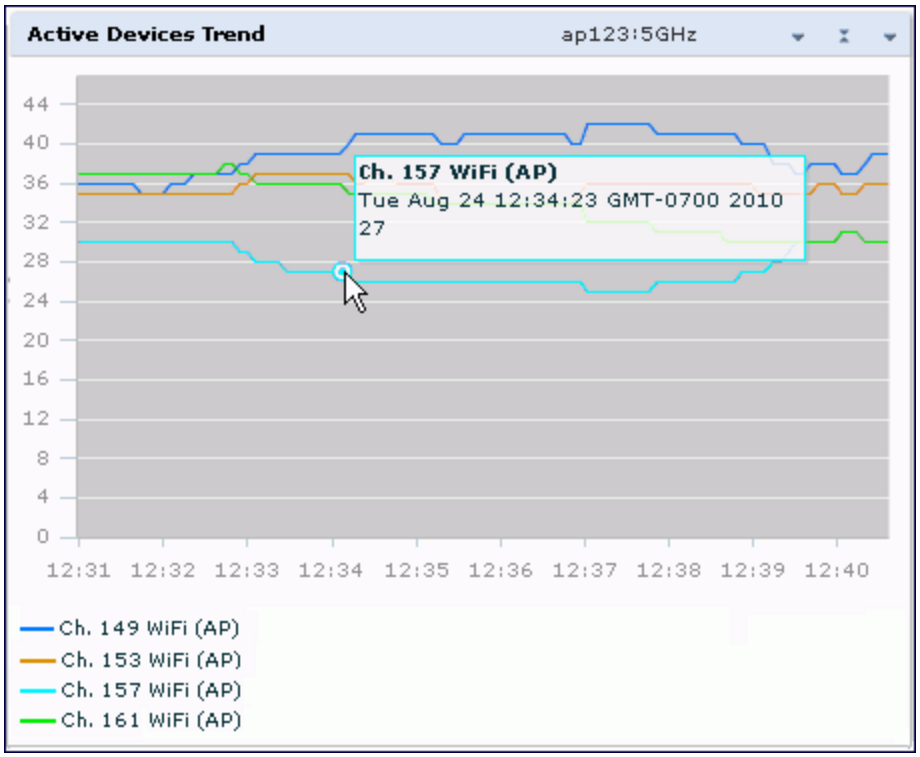
Parameter	Description
	To filter the output of this table to show devices discovered within a specific time, click the icon in the column heading. Select **Any** to display all entries, regardless of when the device was discovered. To display entries for devices discovered within a specific time range: 1. Select the button by the **Less than** drop down list. 2. Click the **Less than** drop-down list and select either **Less than** or **More than** to limit the output of this table to devices discovered earlier or after a specified number of hours or minutes. 3. Enter the number of hours or minutes in the time range you want apply to this filter. 4. Click the **min.** drop down list and select either **min.** or **hrs.** to define the time range in minutes or hours. 5. Click **OK** to save your settings and return to the **Active Devices** table.
Activity Duration	Amount of time that the device has been active. To filter the output of this table to show devices that have been active within a specific time range, click the icon in the column heading. Select **Any** to display all entries, regardless of how long the device has been active. To display entries for devices active for a specific time range: 1. Select the button by the **>** symbol. 2. Click drop-down list with the **>** symbol and select either **>** (greater than), **<** (less than), **<=** (less than or equal to), or **>=** (more than or equal to) to limit the output of this table to devices that have been active for a specified time range. 3. Enter the number of hours or minutes in the time range you want apply to this filter. 4. Click the **min.** drop down list and select either **min.** or **hrs.** to define the time range in minutes or hours. 5. Click**OK** to save your settings and return to the **Active Devices** table.
Channels Affected	Radio channels affected by the device's transmission. The Active Devices table for a spectrum monitor shows entries for all devices by default, regardless of the channels their transmissions may affect. To filter the output of this table to show devices that affect a specific channel or range of channels, click the icon in the column heading. • Select **Any** to display all entries, regardless of the channels that device may affect. • Select **Single Channel**, then enter the channel value to only display devices that affect the specified channel. • Select **Range of Channels**, then enter the lower and upper channels in the channel range to filter the output to show only those devices whose transmissions affect the specified channel range. This option is only available for tables created by spectrum monitors, not hybrid APs. • Select **Specified Channels** to show only those devices whose transmissions affect selected channels. If you choose this option, you can click the **none** checkbox to show only those devices whose transmissions do not affect any other channels, select **all** to show devices whose transmissions affect any channel, or click the checkboxes by individual channel numbers to show only those devices whose output affect those selected channels. This option is only available for tables created by spectrum monitors, not hybrid APs. • Click **OK** to save your settings and return to the **Active Devices** table. **NOTE:** This option is not available for Active Devices tables created by a hybrid AP, because each hybrid AP monitors a single channel only.

Parameter	Description
Device ID	The spectrum monitor or hybrid AP applies a unique device ID per device type to each device it detects on the radio channel. To display the entry for a device that matches a single device ID, click the icon in the column heading and enter the device ID. Click **OK** to save your settings and return to the **Active Devices** table.
Center Frequency (MHz)	Signals from a wireless device can spread beyond the boundaries of an individual 802.11 channel. This table column shows the center frequency for the device's transmission, in megahertz.
Occupied Bandwidth	Channel bandwidth used by the device, in megahertz.

Active Devices Trend

The Active Devices Trend chart is a line chart that shows the numbers of Wi-Fi and non-Wi-Fi devices seen on each radio channel during the displayed time interval. When you hover your mouse over any line in the chart, a tooltip displays the number of active devices for the selected device type. The example in Figure 110 shows that there are 27 active Wi-Fi APs on channel 157 of the 5 GHz radio band.

Figure 110 *Active Devices Trend Graph*



An Active Devices Trend chart created by a hybrid AP displays data for the single channel monitored by that device. For spectrum monitors, the Active Devices Trend chart can display values for up to five different channels and device types. These graphs show the following data by default:

- For SMs on the 2.4 GHz radio band, Wi-Fi APs on channel 1, 6, and 11.
- For SMs on the 5 GHz band, Wi-Fi APs on channel 36, 40, and 44.

Table 137 describes the other values that can be displayed in the Active Devices Trend chart. Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access the Active Devices Trend

configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 137: *Active Devices Trend Options*

Parameter	Description
Band	Radio band displayed in this graph (2.4 GHz or 5 GHz).
Show Trend for Last	Amount of elapsed time for which this chart should display data.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the Channel Numbering drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.
Show lines for these channels	The Active Devices Trend chart can display values for up to five different device types on different channels for a spectrum monitor, or a single device type for a hybrid AP. To choose which type of data each line should represent, click the channel number drop-down list and select a channel within the radio band, then click the device type drop-down list and select one of the following device types. • WiFi (AP) • Microwave *(This option is only available for 2.4 GHz radios)* • Bluetooth *(This option is only available for 2.4 GHz radios)* • Fixed Freq (Others) • Fixed Freq (Cordless Phones) • Fixed Freq (Video) • Fixed Freq (Audio) • Freq Hopper (Others) • Freq Hopper (Cordless Network) • Freq Hopper (Cordless Base) • Freq Hopper Xbox *(This option is only available for 2.4 GHz radios)* • Microwave (Inverter) *(This option is only available for 2.4 GHz radios)* • Generic Interferer Select the checkbox beside each channel and device entry to show that information on the chart, or deselect the checkbox to hide that information. For more information on non-Wi-Fi device types detected by a spectrum monitor, see Working with Non-Wi-Fi Interferers on page 739.

Channel Metrics

This stacked bar chart can show one of three different types of channel metrics; **channel utilization**, **channel availability**, or **channel quality**.

This chart displays channel utilization data by default, showing both the percentage of each monitored channel that is currently being used by Wi-Fi devices, and the percentage of each channel being used by non-Wi-Fi devices and 802.11 adjacent channel interference (ACI).



ACI refers to the interference on a channel created by a transmitter operating in an adjacent channel. A transmitter on a nonadjacent or partially overlapping channel may also cause interference, depending on the transmit power of the interfering transmitter and the distance between the devices. In general, ACI may be caused by a Wi-Fi transmitter or a non-Wi-Fi interferer. However, whenever the term ACI appears in Spectrum Analysis graphs, it refers to the ACI caused by Wi-Fi transmitters. The channel utilization option in the Channel Metrics Chart shows the percentage of the channel utilization due to both ACI and non-Wi-Fi interfering devices. Unlike the ACI shown in the **Interference Power** chart, the ACI shown in this graph indicates the percentage of channel time that is occupied by ACI or unavailable for Wi-Fi communication due to ACI.

The Channel Metrics graph can also show channel availability, the percentage of each channel that is available for use, or display the current relative quality of selected channels in the 2.4 GHz or 5 GHz radio bands. Spectrum monitors can display data for all channels in their selected band. Hybrid APs display data for their one monitored channel only.

In the spectrum analysis feature, channel quality is a relative measure that indicates the ability of the channel to support reliable Wi-Fi communication. Channel quality, which is represented as a percentage in this chart, is a weighted metric derived from key parameters that can affect the communication quality of a wireless channel, including noise, non-Wi-Fi (interferer) utilization and duty-cycles, and certain types of retries. Note that channel quality is not directly related to Wi-Fi channel utilization, as a higher quality channel may or may not be highly used.

When you hover your mouse over any bar in the chart, a tooltip displays the metric value for that individual channel. The example below shows that 61% of channel 3 is being consumed by non-Wi-Fi devices and 802.11 adjacent channel interference.

Figure 111 *Channel Metrics Graph*

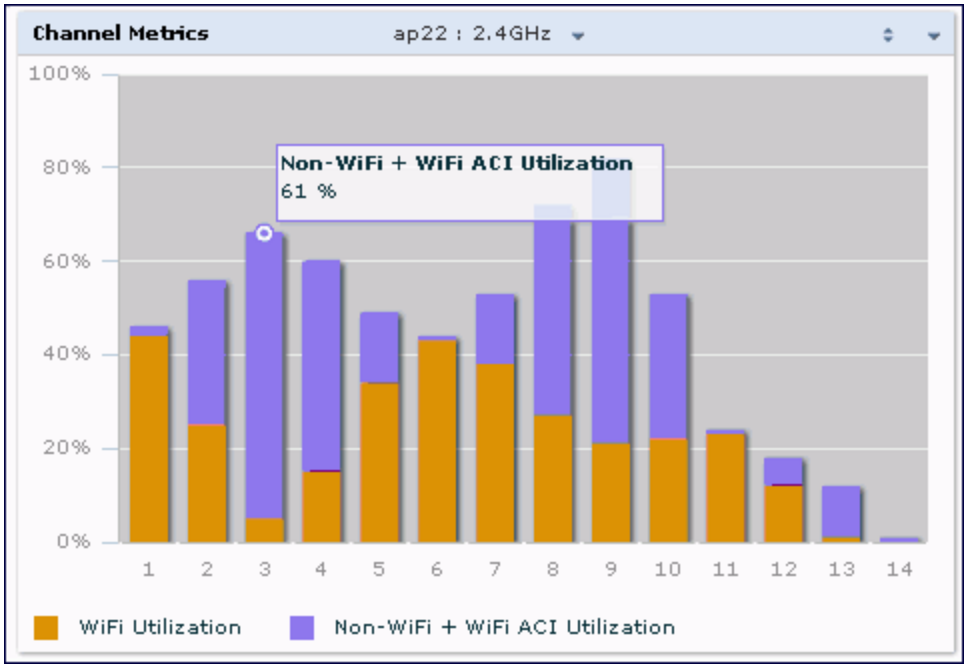


Table 138 describes the parameters that can be displayed in the Channel Metrics graph. Click the down arrow in the upper right corner of this chart then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 138: *Channel Metrics Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz upper**, **5GHz middle** or **5Ghz lower**to display data for that portion of the 5 Ghz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.
Channel Range	For graphs created by spectrum monitors, specify a channel range to determine which channels appear in this graph. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the graph. This graph displays all channels within the spectrum monitor's radio band by default. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.
Display Mode	Select **Channel Quality** to show the relative quality of the channel. Channel Quality is a weighted metric derived from key parameters which include noise, non-Wi-Fi (interferer) utilization and duty-cycles, and certain types of retries. Select **Channel Availability** to show the percentage of the channel that is unused and available for additional Wi-Fi traffic. Select **Channel Utilization** to show both the percentage of the channel that is currently used by Wi-Fi devices, and the percentage of each channel that is being used by non-802.11 devices or 802.11 adjacent channel interference (ACI).

Channel Metrics Trend

By default, this line chart shows the current relative quality of selected channels in the 2.4 GHz or 5 GHz radio bands over a period of time. The Channel Metrics Trend chart can also be configured to display trends for the current availability of selected channels, or the percentage of availability for those channels.Spectrum monitors can display data for up to five different channels. Hybrid APs display data for their one monitored channel only.



For more information on how the spectrum analysis feature determines the quality of a channel, see Channel Metrics on page 719.

When you hover your mouse over any line in the chart, a tooltip displays channel quality or availability data for that individual channel at the selected time.

Figure 112 *Channel Metrics Trend Chart*

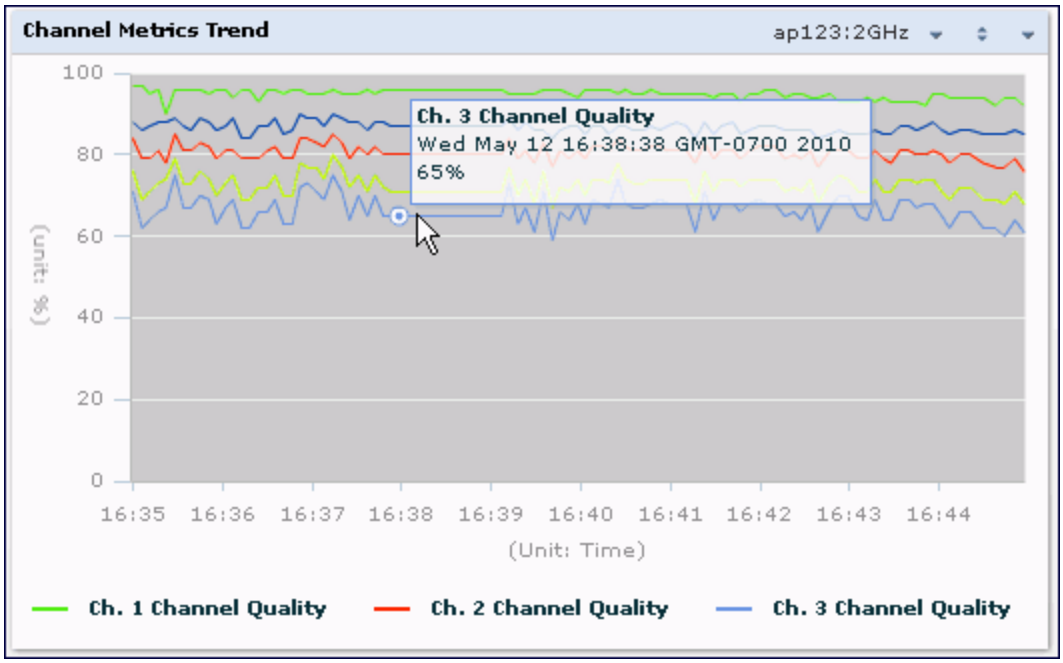


Table 139 describes the other parameters that can be displayed in the Channel Metrics Trend output. Click the down arrow in the upper right corner of this chart then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboard.

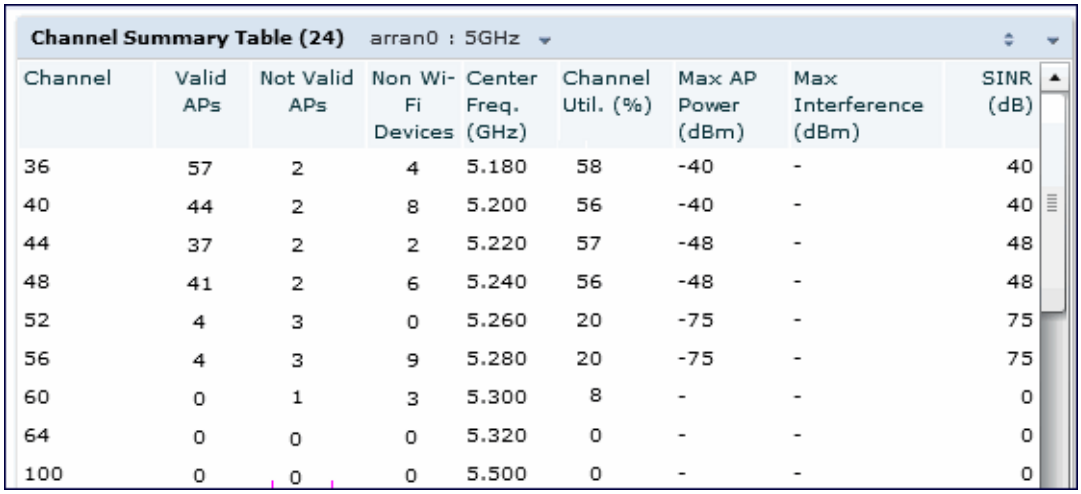
Table 139: *Channel Metrics Trend Options*

Parameter	Description
Band	Radio band displayed in this graph (2.4 GHz or 5 GHz).
Show Trend for Last	The Channel Quality Trend chart shows channel quality or channel availability for the past 10 minutes by default. To view data for a different time range, click the **Show Trend for Last** drop-down list and select one of the following options: • 10 minutes • 30 minutes • 1 hour
Channel numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.
Show Lines for These Channels	The Channel Quality Trend chart for a spectrum monitor can display channel quality, channel availability, or channel utilization values for up to five different channels on the selected radio band. Charts for hybrid APs can display data for the one channel monitored by that hybrid AP radio. To choose which type of data each line should represent on a chart for a spectrum monitor, click the **channel number** drop-down list and select a channel within the radio band, then click the second drop-down list and select either **Channel Quality**, or **Channel Availability**. Select the checkbox beside each channel entry to show that information on the chart, or deselect the checkbox to hide that information.

Channel Summary Table

The channel summary table provides a summarized or aggregated view of key statistics. Spectrum monitors display spectrum analysis data seen on all channels in the selected band, and hybrid APs display data from the one channel they are monitoring. The example in Figure 113 below shows that a spectrum monitor sees 44 Valid APs and 52% channel utilization on channel 40 in the 5GHz radio band.

Figure 113 *Channel Summary Table*



Channel Summary Table (24) arran0 : 5GHz

Channel	Valid APs	Not Valid APs	Non Wi-Fi Devices	Center Freq. (GHz)	Channel Util. (%)	Max AP Power (dBm)	Max Interference (dBm)	SINR (dB)
36	57	2	4	5.180	58	-40	-	40
40	44	2	8	5.200	56	-40	-	40
44	37	2	2	5.220	57	-48	-	48
48	41	2	6	5.240	56	-48	-	48
52	4	3	0	5.260	20	-75	-	75
56	4	3	9	5.280	20	-75	-	75
60	0	1	3	5.300	8	-	-	0
64	0	0	0	5.320	0	-	-	0
100	0	0	0	5.500	0	-	-	0

Spectrum monitor radios using the 5 GHz radio band can display channels using either 20 MHz or 40 MHz channel numbering. Spectrum monitor radios that support 802.11ac can also display 80MHz channels. To toggle between these channel numbering modes, click the down arrow in the upper right corner of the graph titlebar, then click either **Show 20 MHz Channels**, **Show 40 MHz Channels** or **Show 80 MHz Channels.**

Click any of the column headings to sort the information in the table by that column criteria. Make a column wider or narrower by clicking the border of a column heading and dragging the border to a new position.

Table 140 describes the output of the Channel Summary table.

Table 140: *Channel Summary Table Parameters*

Parameter	Description
Channel	Radio channel being monitored by the spectrum monitor or hybrid AP
Valid APs	Number of known APs seen on the network.
Not Valid APs	Number of unknown or invalid APs seen on the network.
Non Wi-Fi Devices	Number of Non-Wi-Fi (interfering) devices detected/classified by the spectrum monitor.
Center Freq. (GHz)	Center frequency of the Wi-Fi signals sent on that radio channel.
Channel Util. (%)	Percentage of the channel currently being used by devices on the network.

Parameter	Description
Max AP Power (dBm)	Signal strength of the AP that has the maximum signal strength on a channel.
Max Interference (dBm)	Signal strength of the non-Wi-Fi device that has the highest signal strength.
SNIR (dB)	The Signal-to-Noise-and-Interference Ratio (SNIR) is the ratio of signal strength to the combined levels of interference and noise on that channel. This value is calculated by determining the maximum noise-floor and interference-signal levels, and then calculating how strong the desired signal is above this maximum.

Device Duty Cycle

The Device Duty Cycle Chart is a stacked bar chart that shows the duty cycle of each device type on a channel. The duty cycle is the percentage of time each device type operates or transmits on that channel. Though Wi-Fi devices do not transmit if there is another Wi-Fi or non-Wi-Fi device active at that time, most non-Wi-Fi devices do not follow such a protocol for transmissions. Because these devices operate independently without regard to any other devices operating on the same channel, the total duty cycle of all device types may add up to more than 100% on a channel. For example, one or more video bridges may be active on a channel, each with a 100% duty cycle. The same channel may have a cordless transmitter with a 10% duty cycle and a microwave oven with a 50% duty cycle. In this example, the Device Duty Cycle chart shows all three device types with their respective duty cycle percentages.



This chart is not available for W-AP68 access points. A hybrid AP on a 20 MHz channel will see 40 MHz Wi-Fi data as non-Wi-Fi data.

Spectrum monitors display spectrum analysis data seen on all channels in the selected band, and hybrid APs display data from the one channel they are monitoring. The example below shows data from a spectrum monitor monitoring all channels in the 2.4 Ghz band.

Figure 114 *Device Duty Cycle*

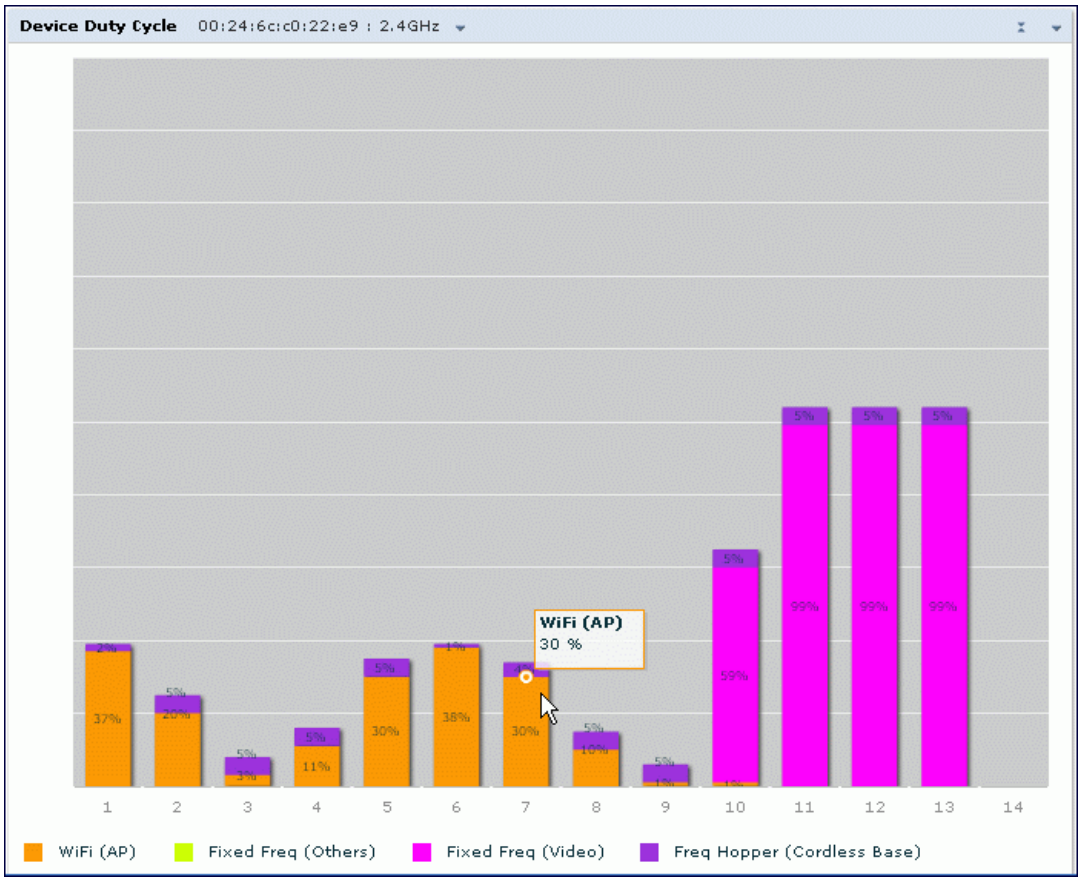


Table 141 describes the parameters you can use to customize the Device Duty Cycle chart. Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 141: *Device Duty Cycle Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz upper**, **5 GHz middle** or **5 Ghz lower** to display data for that portion of the 5 Ghz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80 MHz** option for very-high-throughput channels.
Channel Range	For graphs created by spectrum monitors, specify a channel range to determine which channels appear in this graph. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the graph. This graph displays all channels within the spectrum monitor's radio band by default. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.

Parameter	Description
Show	This graph can display values for up to five different device types on different channels for a spectrum monitor, or a single device type for a hybrid AP monitoring a single channel. To choose which type of data each line should represent, click the channel number drop-down list and select a channel within the radio band, then click the device type drop-down list and select one of the following device types • WiFi (AP) • Microwave *(This option is only available for 2.4 GHz radios)* • Bluetooth *(This option is only available for 2.4 GHz radios)* • Fixed Freq (Others) • Fixed Freq (Cordless Phones) • Fixed Freq (Video) • Fixed Freq (Audio) • Freq Hopper (Others) • Freq Hopper (Cordless Network) • Freq Hopper (Cordless Base) • Freq Hopper Xbox *(This option is only available for 2.4 GHz radios)* • Microwave (Inverter) *(This option is only available for 2.4 GHz radios)* • Generic Interferer **NOTE:** For more information on non-Wi-Fi device types detected by a spectrum monitor, see Working with Non-Wi-Fi Interferers on page 739.

Channel Utilization Trend

The Channel Utilization Trend chart is a line chart that shows the percentage of total utilization on each channel over a time interval. The channel utilization includes the utilization due to Wi-Fi as well as utilization due to non-Wi-Fi interferers and Adjacent Channel Interference (ACI).



For additional information on how the spectrum analysis feature measures ACI, see Channel Metrics on page 719.

This graph can show data recorded for the last ten, thirty, or sixty minutes. Spectrum monitors display spectrum analysis data seen on all channels in the selected band, and hybrid APs display data from the one channel they are monitoring. When you hover your mouse over any line in the chart, a tooltip shows the percentage of the channel being utilized at the specified time. The example in Figure 115 shows that channel 1 was 70% used at the selected time in the chart.

Figure 115 *Channel Utilization Trend*

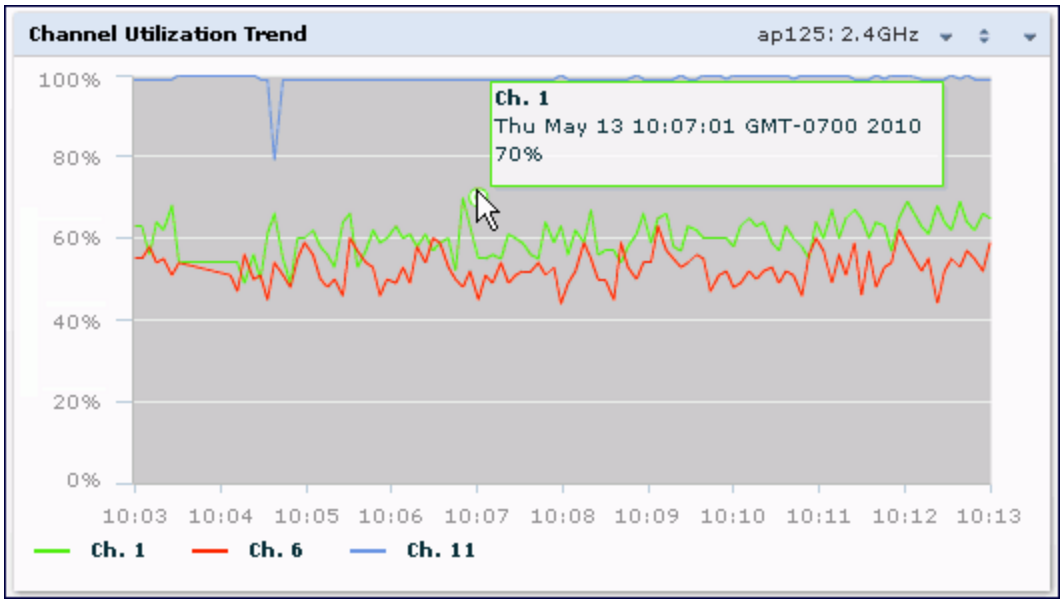


Table 142 describes the parameters you can use to customize the Channel Utilization Trend chart. Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 142: *Channel Utilization Trend Options*

Parameter	Description
Intervals	The Channel Utilization Trend chart shows channel quality or channel availability for the past 10 minutes by default. To view data for a different time range, click the **Intervals** drop-down list and select one of the following options: • 10 minutes • 30 minutes • 1 hour
Band	Radio band displayed in this graph (2.4 GHz or 5 GHz).
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.
Show	To select individual channels you want to display on this chart, click the checkbox by a channel entry, then click the **channel** drop-down list to select the channel to display. To hide a channel, uncheck the checkbox by that channel number.

Devices vs Channel

This stacked bar chart shows the current number of devices using each channel in the radio's frequency band. This chart can show separate per-channel statistics for the numbers of Wi-Fi devices, cordless phones, bluetooth devices, microwaves, and other non-Wi-Fi devices.

If a device affects more than one channel, it is recorded as a device on all channels it affects. For example, if a 20Mhz Wi-Fi AP has a center frequency of 2437 Mhz (channel 6) it is counted as a device on channels 3-9 because it affects all those channels. Similarly, if a channel-hopping device uses all channels within a frequency band, it is counted as a device on all channels in that band.

When you hover the mouse over any part of the chart, a tooltip shows the numbers of the device type currently using that channel. The example in Figure 116 shows that the spectrum monitor can detect 42 APs on channel 5.

Figure 116 *Devices vs Channel*

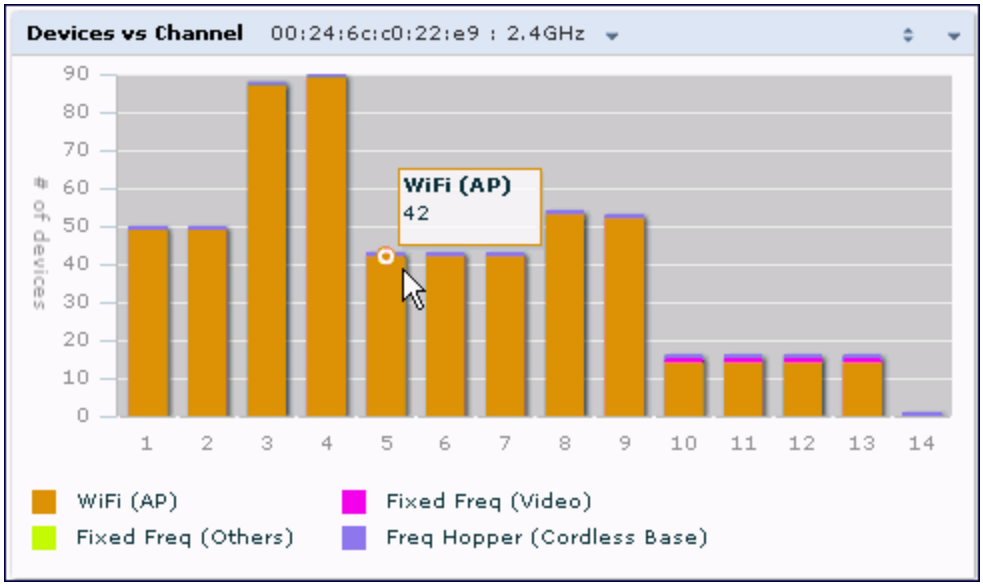


Table 143 describes the parameters you can use to customize the Devices vs Channel chart. Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 143: *Devices vs Channel Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz upper**, **5GHz middle** or **5Ghz lower** to display data for that portion of the 5 Ghz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.
Channel Range	For graphs created by spectrum monitors, specify a channel range to determine which channels appear in this graph. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the graph.

Parameter	Description
	This graph displays all channels within the spectrum monitor's radio band by default. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.
Show	This graph can show data for up to five different device types. To show how many devices of a specific type are sending a signal on the selected channel range, click the **show** checkbox by that device, then click the device drop-down list and select one of the following device types. • WiFi (AP) • Microwave *(This option is only available for 2.4 GHz radios)* • Bluetooth *(This option is only available for 2.4 GHz radios)* • Fixed Freq (Others) • Fixed Freq (Cordless Phones) • Fixed Freq (Video) • Fixed Freq (Audio) • Freq Hopper (Others) • Freq Hopper (Cordless Network) • Freq Hopper (Cordless Base) • Freq Hopper Xbox *(This option is only available for 2.4 GHz radios)* • Microwave (Inverter) *(This option is only available for 2.4 GHz radios)* • Generic Interferer **NOTE:** For more information on non-Wi-Fi device types detected by a spectrum monitor, see Working with Non-Wi-Fi Interferers on page 739.

FFT Duty Cycle

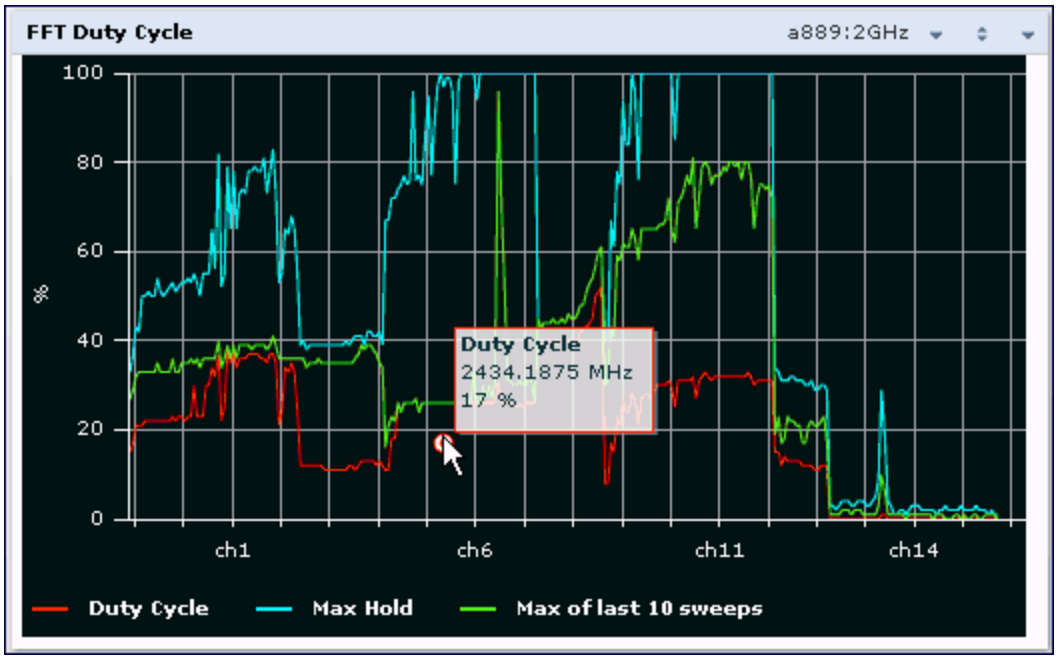
The FFT Duty Cycle chart is a line chart that shows the duty cycle for each frequency bin. The width of the each frequency bin depends on the resolution bandwidth of the spectrum monitor. The spectrum analysis feature considers a frequency bin to be used if the detected power in that bin is at least 20 dB higher than the nominal noise floor on that channel. The FFT Duty Cycle provides a more granular view of the duty cycle per bin as opposed to the aggregated channel utilization reported in the Channel Metrics chart.



This chart is not available for W-AP68 access points. A hybrid AP on a 20 MHz channel will see 40 MHz Wi-Fi data as non-Wi-Fi data.

This chart can show the duty cycle over the last second, the maximum FFT duty cycle measured for all samples taken over the last N sweeps, and the greatest FFT duty cycle recorded since the chart was last reset.

Figure 117 *FFT Duty Cycle*



This chart shows the current duty cycle for devices on all channels being monitored by the spectrum monitor radio by default. Table 144 describes the other optional parameters you can use to customize the FFT Duty Cycle table. Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 144: *FFT Duty Cycle Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz upper**, **5GHz middle** or **5Ghz lower** to display data for that portion of the 5 Ghz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.
X-Axis	Select either **Channel** or **Frequency** to show the duty cycle for a range of channels or frequencies.
Channel Range	If you selected **Channel** in the **X-Axis** parameter, you must also specify a channel range to determine which channels appear in the x-axis of this chart. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the chart. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.

Parameter	Description
Center Frequency	If you selected **Frequency** in the **X-Axis** parameter, enter the frequency, in MHz, that you want to appear in the center of the x-axis of this chart.
Span	If you selected **Frequency** in the **X-Axis** parameter, specify the size of the range of frequencies around the selected center frequency. If you set a frequency span of 100 MHz, for example, the chart shows the FFT duty cycle for a range of frequencies from 50MHz lower to 50 MHz higher than the selected center.
Show	Select a checkbox to display that information on the FFT Duty Cycle chart. • **Duty Cycle:** The percentage of duty cycle the channel or frequency was actively used. • **Max Hold:** The maximum recorded percentage of active duty cycles for the channel frequency since the chart was last reset. To clear this setting, click the down arrow at the end of the title bar for this graph and select **Reset MaxHold**. • **Max of last sweeps:** This chart shows the maximum percentage of active duty cycles for the channel of frequency recorded during the last 10 sweeps, by default. To change the number of sweeps used to determine this value, enter a number from 2 to 20, inclusive. To clear this setting, click the down arrow at the end of the title bar for this graph and select **Reset MaxNSweep**.

Interference Power

The Interference Power chart displays various power levels of interest, including the Wi-Fi AP with maximum signal strength, noise, and interferer types with maximum signal strength. The ACI displayed in the Interference Power Chart is the ACI power level based on the signal strength(s) of the Wi-Fi APs on adjacent channels. A higher ACI value in Interference Power Chart does not necessarily mean higher interference, because the AP that is contributing to the maximum ACI may or may not be very actively transmitting data to other clients at all times. The ACI power levels are derived from the signal strength of the beacons.

This chart displays the noise floor of each selected channel in dBm. The noise floor of a channel depends on the noise figure of the RF components used in the radio, temperature, presence of certain types of interferers or noise, and the width of the channel. For example, in a clean RF environment, a 20 MHz channel has a noise floor around -95 dBm and a 40 MHz channel has a noise floor around -92 dBm. Certain types of fixed-frequency continuous transmitters such as video bridges, fixed-frequency phones, and wireless cameras typically elevate the noise floor seen by the spectrum monitor. Other interferers such as frequency-hopping phones, Bluetooth, and Xbox may not affect the noise floor of the radio. A Wi-Fi radio can only reliably decode Wi-Fi signals that are a certain dB above the noise floor. Therefore estimating and understanding the actual noise floor of the radio is critical to understanding the reliability of the RF environment.

The chart also includes information about the AP on each channel with the highest power level. You can hover your mouse over an AP on the chart to view the AP's name, SSID, and current power level. The example below shows that the AP with the maximum power on channel 157 has the SSID **qa-ss**, and a power level of -55dBm.

Figure 118 *Interference Power*

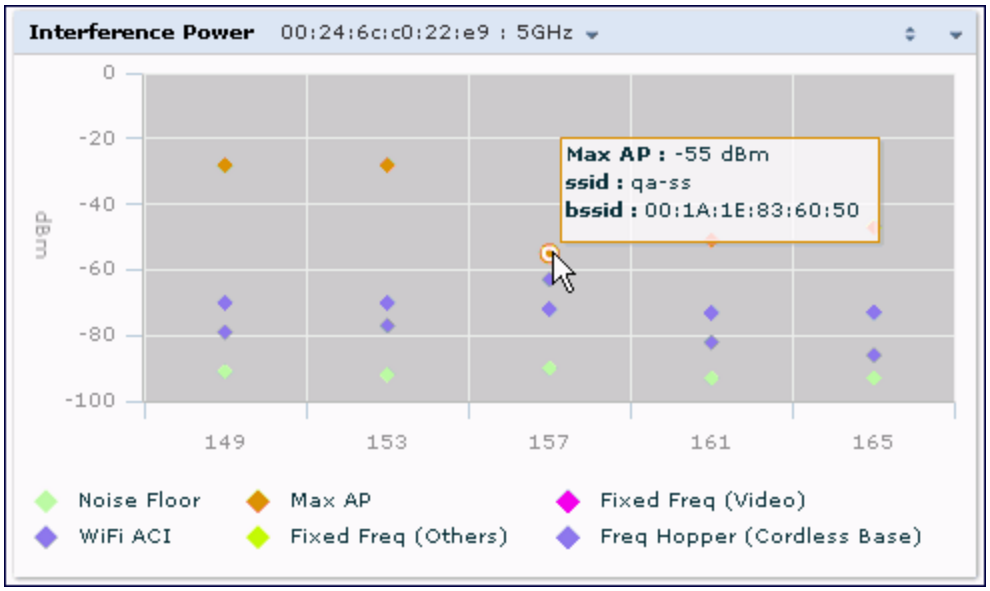


Table 145 describes the other optional parameters you can use to customize the interference power chart. Click the down arrow in the upper right corner of this chart then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 145: *Interference Power Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz upper**, **5GHz middle** or **5Ghz lower** to display data for that portion of the 5 Ghz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.

Parameter	Description
Show	By default, this chart displays data for the current noise floor, adjacent channel interference (ACI), and the maximum AP power level for each channel. To display interference power levels form other devices, click the show checkbox then click the **show** drop-down list and select one of the following device types. • Microwave *(This option is only available for 2.4 GHz radios)* • Bluetooth *(This option is only available for 2.4 GHz radios)* • Fixed Freq (Others) • Fixed Freq (Cordless Phones) • Fixed Freq (Video) • Fixed Freq (Audio) • Freq Hopper (Others) • Freq Hopper (Cordless Network) • Freq Hopper (Cordless Base) • Freq Hopper Xbox *(This option is only available for 2.4 GHz radios)* • Microwave (Inverter) *(This option is only available for 2.4 GHz radios)* • Generic Interferer For more information on non-Wi-Fi device types detected by a spectrum monitor, see Working with Non-Wi-Fi Interferers on page 739.
Channel Range	For graphs created by spectrum monitors, specify a channel range to determine which channels appear in this graph. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the graph. This graph displays all channels within the spectrum monitor's radio band by default. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.

Quality Spectrogram

This plot shows the channel quality statistics for selected range of channels or frequencies. This chart can also be configured to show channel availability, the percentage of each channel that is unused and available for additional traffic.

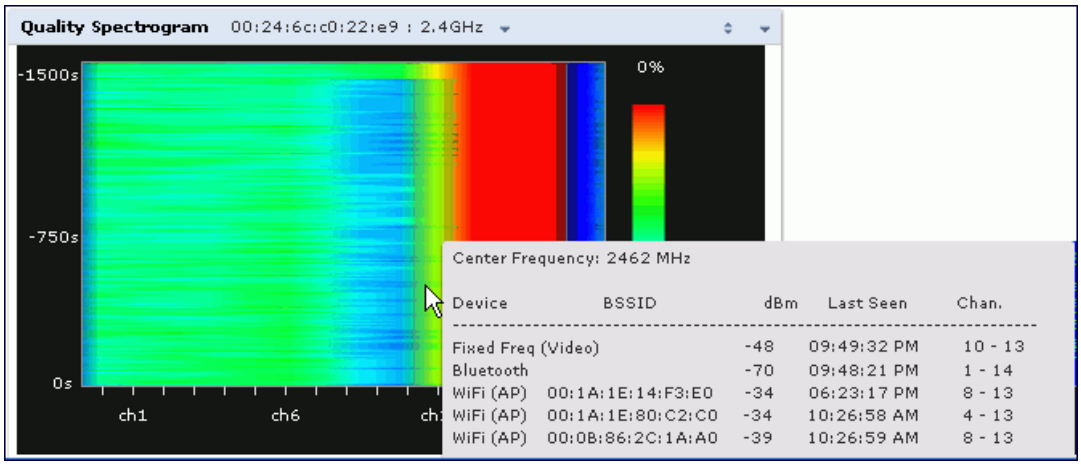
Channel Quality is a weighted metric derived from key parameters which include noise, non-Wi-Fi (interferer) utilization and duty-cycles and certain types of retries. Quality levels are indicated by a range of colors between dark blue, which represents a higher channel quality, and red, which represents a lower channel quality. Channel availability is indicated by a range of colors between dark blue, which represents 100% channel availability, and red, which represents 0% availability.



For additional information on interpreting a Dell Spectrogram plot, see Swept Spectrogram on page 736.

The Spectrum Analysis Quality Spectrogram chart measures channel data each second, so after every 5-second sweep, the newest data appears as a thin colored line on the bottom of the chart. Older data is pushed up higher on the chart until it reaches the top of the spectrogram and ages out. The example below shows the Dell Quality Spectrogram chart after it has recorded over 1,500 seconds of FFT data.

Figure 119 *Quality Spectrogram*



When you hover your mouse over any part of the spectrogram, a tooltip shows the devices the spectrum monitor detected on that frequency, the BSSID of the device (if applicable), the power level of the device in dBm, the time the device was last seen by the spectrum monitor, and the channels affected by the device.

The following table describes the other optional parameters you can use to customize the Quality Spectrogram. Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 146: *Quality Spectrogram Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz upper**, **5GHz middle** or **5Ghz lower** to display data for that portion of the 5 Ghz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.
Channel Range	Specify a channel range to determine which channels appear in the x-axis of this chart. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the chart. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.

Real-Time FFT

The Real-time FFT chart displays the instantaneous Fast Fourier Transform (FFT) signature of the RF signal seen by the radio. The Fast Fourier Transform (FFT) converts an RF signal from time domain to frequency domain. The frequency domain representation divides RF signals into discrete frequency bins; small frequency ranges whose width depends on the resolution bandwidth of the spectrum monitor (that is, how many Hz are represented by a single signal strength value). Each frequency bin has a corresponding signal strength value.

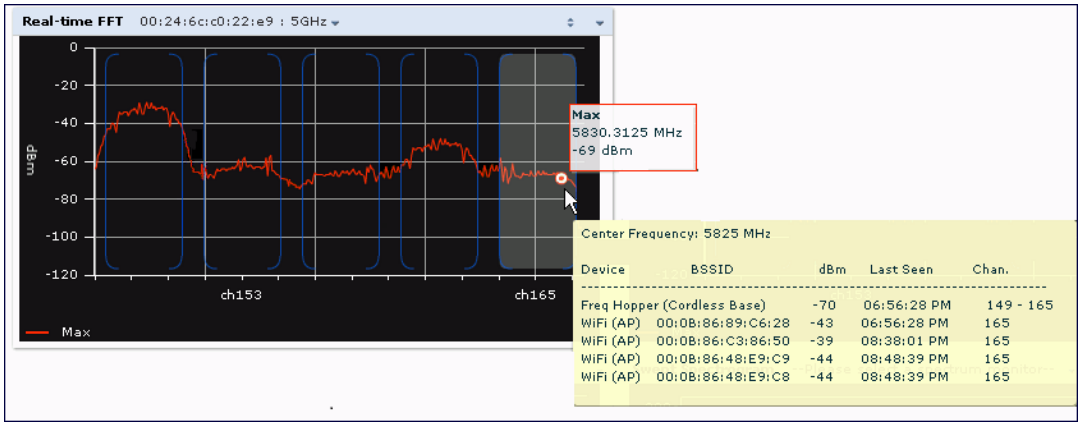
Because there may be a large number of FFT signatures received by the radio every second, an algorithm selects one FFT sample to display in the Real-time FFT chart every second.



This chart is not available for W-AP68 access points. A hybrid AP on a 20 MHz channel will see 40 MHz Wi-Fi data as non-Wi-Fi data.

This chart can show an average for all samples taken over the last second, the maximum FFT power measured for all samples taken over ten channel sweeps, and the greatest FFT power recorded since the chart was last reset. When you hover your mouse over any line, a tooltip shows the power level and channel or frequency level represented by that point in the graph. When you hover your mouse over a frequency level (within the blue brackets on the graph), a tooltip shows the types of devices seen on that frequency, and each device's BSSID, power level, channels affected and the time the device was last seen by the spectrum monitor.

Figure 120 *Real-TIme FFT*



This chart shows the maximum power level recorded for any device on all channels or frequencies monitored by the spectrum monitor radio by default.

Table 147 describes the other parameters you can use to customize the Real-time FFT chart. Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 147: *Real-Time FFT Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz upper**, **5GHz middle** or **5Ghz lower**to display data for that portion of the 5 Ghz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.

Parameter	Description
X-Axis	Select either **Channel** or **Frequency** to show FFT power for a range of channels or frequencies. If you select **Frequency**, you must select the radio frequency on which this chart should center, and determine the span of frequencies for the graph.
Channel Range	If you selected **Channel** in the **X-Axis** parameter, you must also specify a channel range to determine which channels appear in the X-axis of this chart. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the chart. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.
Center Frequency	If you selected **Frequency** in the **X-Axis** parameter, enter the frequency, in MHz, that you want to appear in the center of the x-axis of this chart.
Span	If you selected **Frequency** in the **X-Axis** parameter, specify the size of the range of frequencies around the selected center frequency. If you set a frequency span of 100 MHz, for example, the chart shows the FFT duty cycle for a range of frequencies from 50MHz lower to 50 MHz higher than the selected center.
Y-axis	Select the range of power levels, in -dBm, to appear in the y-axis of this chart. Enter the lower value in the right field, and the higher value in the left field.
Show	Select the checkbox by the following items to display that information on the FFT Power chart. • **Average:** the average power level of all samples recorded during the last 10 sweeps. • **Maxthe** The highest power recorded during the last 10 channel sweeps. • **Max Hold:** the highest maximum power level recorded since the chart data was reset. To clear this setting, click the down arrow at the end of the title bar for this graph and select **Clear Max Hold**.

Swept Spectrogram

A spectrogram is a chart that shows how the density of the quantity being plotted varies with time. The spectrum analysis Swept Spectrogram chart plots real-time FFT Maximums, real-time FFT Averages, or the FFT Duty Cycle. In this swept spectrogram, the x-axis represents frequency or channel and the y-axis represents time. Each line in the swept spectrogram corresponds to the data displayed in the Real-Time FFT or FFT Duty Cycle chart.

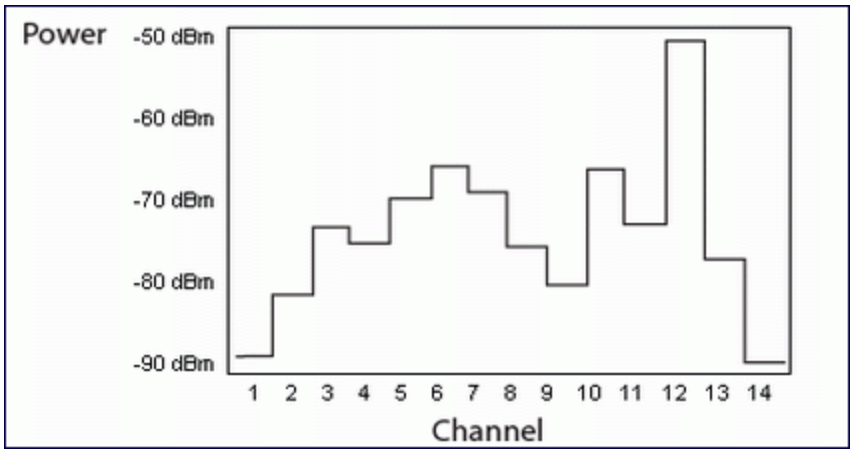


This chart is not available for W-AP68 access points. A hybrid AP on a 20 MHz channel will see 40 MHz Wi-Fi data as non-Wi-Fi data.

The power or duty cycle values recorded in each sweep are mapped to a range of colors. In the average or maximum FFT power Swept Spectrogram charts, the signal strength levels are indicated by a range of colors between dark blue, which represents -90 dBm, and red, which represents a higher -50 dBm. The duty cycle Swept Spectrogram chart shows the percentage of the time tick interval that the selected channel or frequency was broadcasting a signal. These percentages are indicated by a range of colors between dark blue, which represents a duty cycle of 0% percent, and red, which represents a duty cycle of 100%.
A spectrogram plot is a complex chart that can display a lot of information. If you are not familiar with these types of charts, they may be difficult to interpret. The following illustrations can help explain how FFT power data is rendered in a spectrogram format.

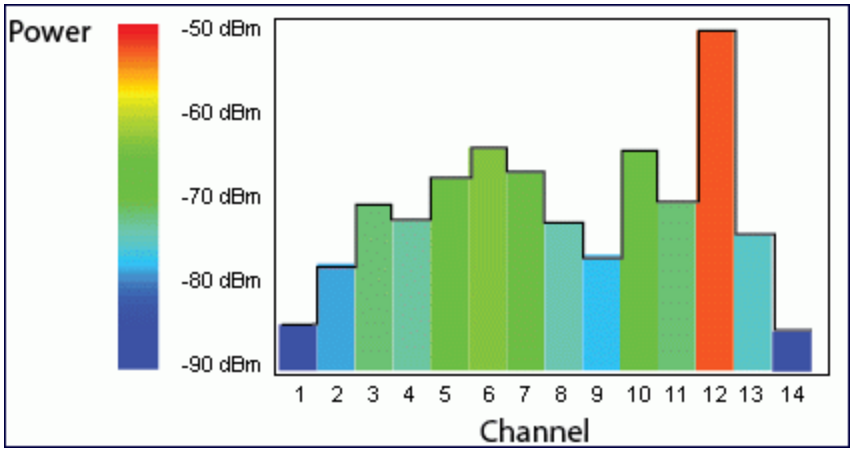
The example in Figure 121 shows how an FFT Power chart could appear if a single data measurement was plotted as a simple line graph.

Figure 121 *Simple Line Graph of FFT Power Data*



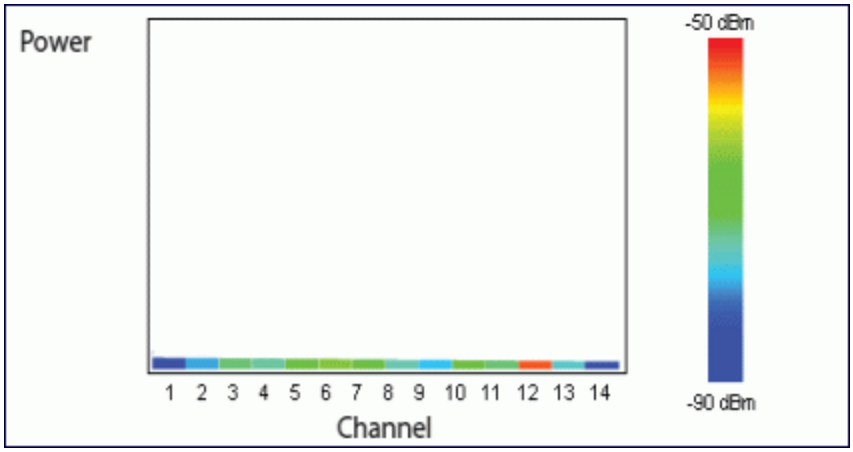
Now, suppose that each channel's FFT power level was also represented by a color that corresponded to that specific FFT power level. In the example below, channel 12 has a FFT power level of -50 dBm, represented by the color red. Channel 1 has a FFT power level of -85 dBm, represented by dark blue.

Figure 122 *FFT Power Line Graph with Color*



If the graph was then flattened so each channel's FFT power for that single1-second sweep was represented only by a color (and not by a value on the y-axis), the graph could then appear as follows:

Figure 123 *FFT Power Spectrogram Sample*



The spectrum analysis Swept Spectrogram measures FFT power levels or duty cycle data each second, so after every 1-second sweep, the newest data appears as a thin colored line on the bottom of the chart. Older data is pushed up higher on the chart until it reaches the top of the spectrogram and ages out. The example below shows the Swept Spectrogram chart after it has recorded over 300 seconds of FFT data.

Figure 124 *Swept Spectrogram*

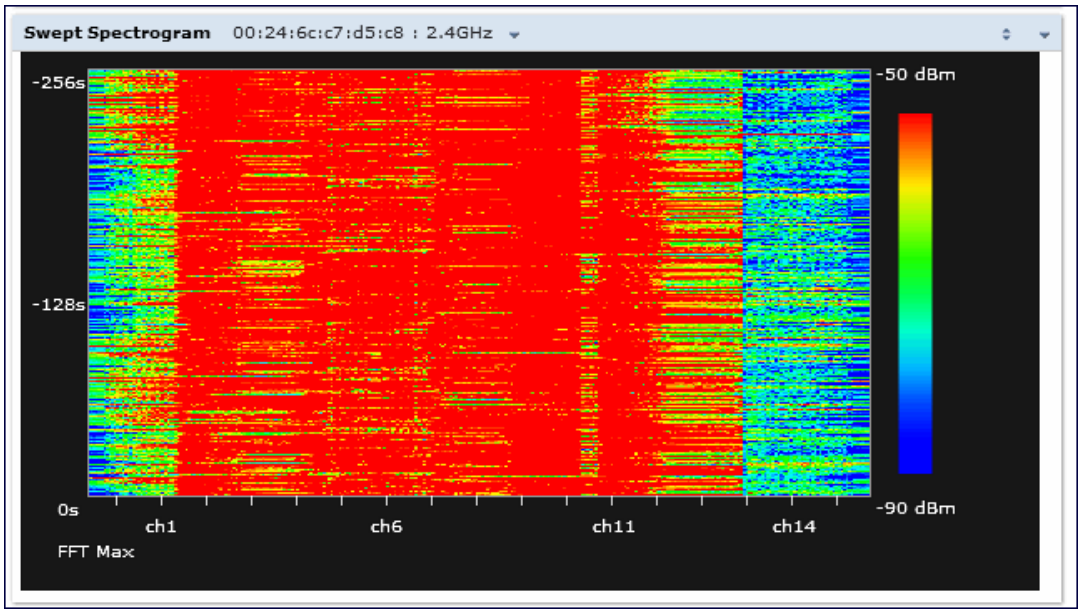


Table 148 describes the parameters you can use to customize the Swept Spectrogram chart. Click the down arrow in the upper right corner of this chart, then click the **Options** menu to access these configuration settings. Once you have configured the desired parameters, click **OK** at the bottom of the **Options** menu to save your settings and return to the spectrum dashboards.

Table 148: *Swept Spectrogram Options*

Parameter	Description
Band	Radio band displayed in this graph. For spectrum monitor radios using the 5 GHz radio band, click the **Band** drop-down list and select **5 GHz upper, 5GHz middle**, or **5Ghz lower** to display data for that portion of the 5 Ghz radio band. This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band.
Channel Numbering	This parameter is not configurable for graphs created by hybrid APs or spectrum monitor radios that use the 2.4 GHz radio band. A hybrid AP on a 20 MHz channel sees 40 MHz Wi-Fi data as non-Wi-Fi data. For spectrum monitors using the 5 GHz radio band, click the **Channel Numbering** drop-down list and select either **20 MHz** or **40 MHz** channel numbering to identify a channel numbering scheme for the graph. Graphs for AP radios that support 802.11ac include an additional **80MHz** option for very-high-throughput channels.
X-Axis	Select either **Channel** or **Frequency** to show FFT power or duty cycles for a range of channels or frequencies. If you select **Frequency**, you must select the radio frequency on which this chart should center, and determine the span of frequencies for the graph.

Parameter	Description
Channel Range	If you selected **Channel** in the **X-Axis** parameter, you must also specify a channel range to determine which channels appear in the x-axis of this chart. Click the first drop-down list to select the lowest channel in the range, then click the second drop-down list to select the highest channel to appear in the chart. **NOTE:** This parameter is not configurable for graphs created by hybrid APs.
Center Frequency	If you selected **Frequency** in the **X-Axis** parameter, enter the frequency, in MHz, that you want to appear in the center of the x-axis of this chart.
Span	If you selected **Frequency** in the **X-Axis** parameter, specify the size of the range of frequencies around the selected center frequency. If you set a frequency span of 100 MHz, for example, the chart shows the swept spectrogram for a range of frequencies from 50MHz lower to 50 MHz higher than the selected center.
Color-Map Range	*If this chart is configured to show average or maximum FFT values,* the default color range on this chart represents values from -50dBm (red) to -90dBm (blue). If you would like the color range on this chart to represent a different range of FFT power levels, enter this range in the **from** and **to** entry blanks. For example, if you defined a color-map range from -60 to -80, then any FFT power level at or above -60 dBm appears as red, and any FFT power level at or below -80 appears blue. Only the channel or frequency qualities between -60 dBm and -80 dBm would be represented by gradiented colors within the color range. *If this chart is configured to show the FFT duty cycle*, the default color range on this chart represents duty cycles from 0% (red) to 100% (blue). If you would like the color range on this chart to represent a different range of FFT duty cycle percentages, enter this range in the **from** and **to** entry blanks. For example, if you defined a color-map range from 25 to 75, then any FFT duty cycle at or below 25% appears as red, and any FFT duty cycle at or below 75% appears blue. Only the duty cycle levels between 25% and 75% would be represented by gradiented colors within the color range. **NOTE:** If your swept spectrogram is showing a single color only, you may need to increase the color map range to display a greater range of values.
Show	Select **FFT Avg**, **FFT Max** or **FFT Duty Cycle** to select the type of data you want to appear in this chart.

Working with Non-Wi-Fi Interferers

The following table describes each type of non-Wi-Fi interferer detected by the spectrum analysis feature. These devices appear in the following charts:

- Active Devices
- Active Devices Table
- Active Devices Trend
- Device Duty Cycle
- Device vs Channel
- Interference Power

Table 149: *Non-Wi-Fi Interferer Types*

Non-Wi-Fi Interferer	Description
Bluetooth	Any device that uses the Bluetooth protocol to communicate in the 2.4 GHz band is classified as a *Bluetooth* device. Bluetooth uses a frequency hopping protocol.
Fixed Frequency (Audio)	Some audio devices such as wireless speakers and microphones also use fixed frequency to continuously transmit audio. These devices are classified as *Fixed Frequency (Audio)*.
Fixed Frequency (Cordless Phones)	Some cordless phones use a fixed frequency to transmit data (much like the fixed frequency video devices). These devices are classified as *Fixed Frequency (Cordless Phones)*.
Fixed Frequency (Video)	Video transmitters that continuously transmit video on a single frequency are classified as *Fixed Frequency (Video*). These devices typically have close to a 100% duty cycle. These types of devices may be used for video surveillance, TV or other video distribution, and similar applications.
Fixed Frequency (Other)	All other fixed frequency devices that do not fall into one of the above categories are classified as *Fixed Frequency (Other*). Note that the RF signatures of the fixed frequency audio, video and cordless phone devices are very similar, and that some of these devices may be occasionally classified as Fixed Frequency (Other).
Frequency Hopper (Cordless Base)	Frequency hopping cordless phone base units transmit periodic beacon-like frames at all times. When the handsets are not transmitting (i.e., no active phone calls), the cordless base is classified as *Frequency Hopper (Cordless Base*).
Frequency Hopper (Cordless Network)	When there is an active phone call and one or more handsets are part of the phone conversation, the device is classified as *Frequency Hopper (Cordless Network*). Cordless phones may operate in 2.4 GHz or 5 GHz bands. Some phones use both 2.4 GHz and 5 GHz bands (for example, 5 GHz for Base-to-handset and 2.4 GHz for Handset-to-base). These phones may be classified as unique Frequency Hopper devices on both bands.
Frequency Hopper (Xbox)	The Microsoft Xbox device uses a frequency hopping protocol in the 2.4 GHz band. These devices are classified as *Frequency Hopper (Xbox*).
Frequency Hopper (Other)	When the classifier detects a frequency hopper that does not fall into one of the above categories, it is classified as *Frequency Hopper (Other*). Some examples include IEEE 802.11 FHSS devices, game consoles, and cordless/hands-free devices that do not use one of the known cordless phone protocols.

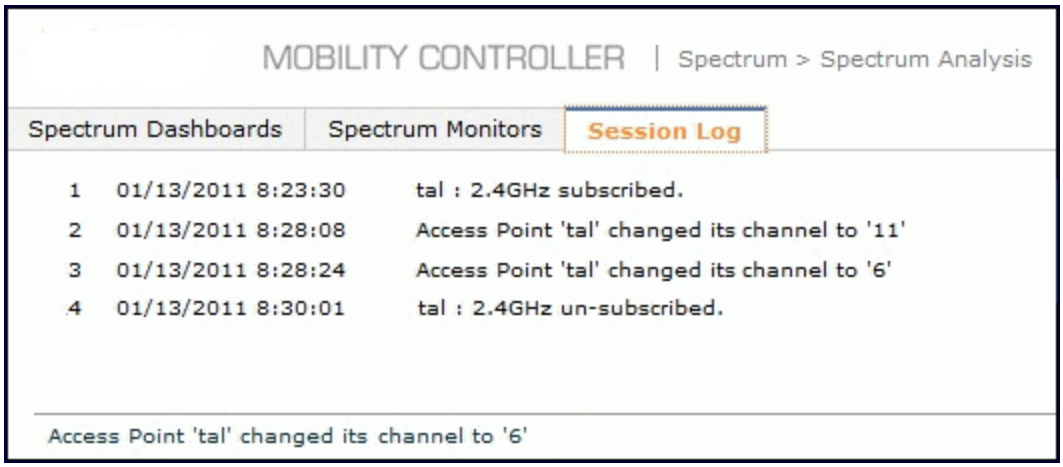
Non-Wi-Fi Interferer	Description
Microwave	Common residential microwave ovens with a single magnetron are classified as a *Microwave*. These types of microwave ovens may be used in cafeterias, break rooms, dormitories and similar environments. Some industrial, healthcare or manufacturing environments may also have other equipment that behave like a microwave and may also be classified as a Microwave device.
Microwave (Inverter)	Some newer-model microwave ovens have inverter technology to control the power output and may have a duty cycle close to 100%. These microwave ovens are classified as *Microwave (Inverter)*. Dual-magnetron industrial microwave ovens with higher duty cycle may also be classified as Microwave (Inverter). As in the Microwave category described above, there may be other equipment that behave like inverter microwaves in some industrial, healthcare or manufacturing environments. Those devices may also be classified as Microwave (Inverter).
Generic Interferer	Any non-frequency hopping device that does not fall into one of the other categories described in this table is classified as a *Generic Interferer*. For example, a Microwave-like device that does not operate in the known operating frequencies used by the Microwave ovens may be classified as a Generic Interferer. Similarly, wide-band interfering devices may be classified as Generic Interferers.

Understanding the Spectrum Analysis Session Log

The spectrum analysis **Session Log** tab displays times the spectrum monitors and hybrid APs connected to or disconnected from the spectrum client during the current browser session. This tab also shows changes in a hybrid AP's scanning channel caused by changes to the hybrid AP's 802.11a or 802.11g radio profile or automatic channel changes by the DFS or ARM features. The latest entry in the session log is also displayed in a footer at the bottom of the Spectrum Monitors and Spectrum Dashboard window. When you close the browser and end your spectrum analysis session, the session log is cleared.

The example in Figure 125 shows that a 2.4 GHz radio on hybrid AP was connected to the spectrum analysis client, its channel changed twice, then was disconnected from the spectrum client.

Figure 125 *Spectrum Analysis Session Logs*



Viewing Spectrum Analysis Data

You can use the command-line interface to view spectrum analysis data from any spectrum monitor, even if that spectrum monitor is currently sending data to another spectrum monitor client's WebUI.

Table 150 shows the commands that display spectrum analysis data in the CLI interface.

Table 150: *Spectrum Analysis CLI Commands*

Command	Description
`show ap spectrum ap-list`	Shows spectrum data seen by an access point that has been converted to a spectrum monitor.
`show ap spectrum channel-metrics`	Shows channel utilization information for a 802.11a or 802.11g radio band, as seen by a spectrum monitor
`show ap spectrum channel-summary`	Displays a summary of the 802.11a or 802.11g channels seen by a spectrum monitor.
`show ap spectrum client-list`	Shows details for Wi-Fi clients seen by a specified spectrum monitor.
`show ap spectrum debug`	Sub-commands under this command save spectrum analysis channel information to a file on the controller.
`show ap spectrum device-duty-cycle`	Shows the current duty cycle for devices on all channels being monitored by the spectrum monitor radio.
`show ap spectrum device-history`	Displays spectrum analysis history for non-interfering devices.
`show ap spectrum device-list`	Shows summary table and channel information for non-Wi-Fi devices currently seen by the spectrum monitor.
`show ap spectrum device-log`	Shows a time log of add and delete events for non-Wi-Fi devices.
`show ap spectrum device-summary`	Shows the numbers of Wi-Fi and non-Wi-Fi device types on each channel monitored by a spectrum monitor.
`show ap spectrum interference-power`	Shows the interference power detected by a 802.11a or 80211g radio on a spectrum monitor.
`show ap spectrum monitors`	Shows a list of APs currently configured as spectrum monitors.
`show ap spectrum technical-support`	Saves spectrum data for later analysis by your Dell technical support representative.

Recording Spectrum Analysis Data

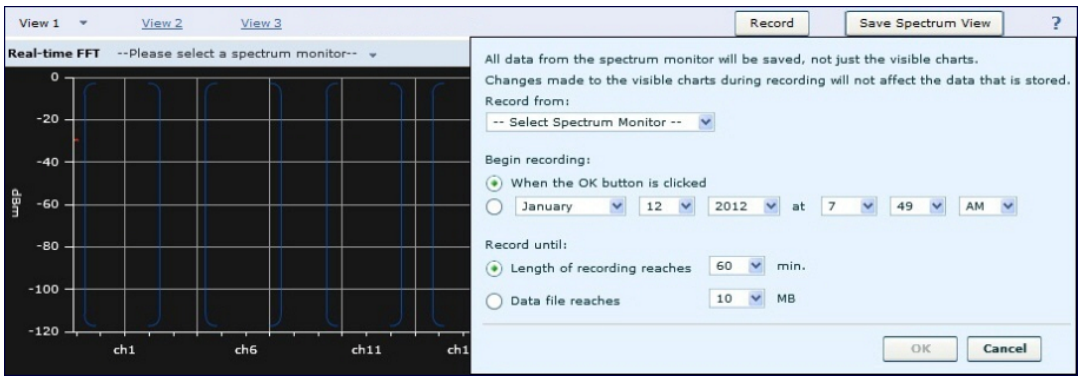
The spectrum analysis tool allows you to record up to 60 continuous minutes (or up to 10 Mb) of spectrum analysis data. By default, each spectrum analysis recording displays data for the Real-Time FFT, FFT Duty Cycle, Interference Power and Swept Spectrogram charts, however, you can view recorded device data for any the spectrum analysis charts supported by that spectrum monitor radio. Configurable recording settings allow you to start a recording session immediately, or schedule a recording to begin at a later date and time. Each recording can be scheduled to end after a selected amount of time has passed, or continue on until the recorded data file reaches a specified size. You can save the file to your spectrum monitor client, then play back that data at a later time.

Creating a Spectrum Analysis Record

To record spectrum analysis data for later analysis:

1. Navigate to the **Monitoring > Spectrum Analysis > Spectrum Dashboards** window.
2. Click **Record** at the top of the window. The **New Recording** popup window appears.
3. Click the **Record From** link, and select the spectrum monitor whose data you want to record.
4. Next, decide whether you want the recording to start immediately, or at a later scheduled time. If you want the recording to start immediately, select **When the OK button is clicked**. To schedule a different starting time for the recording, click the date and time drop-down lists to select a starting month, day, year and time.
5. The recording continues until either the specified amount of time has passed, or until the recording files reaches a selected size. Click the **Length of recording reaches** drop down list and select the amount of time the recording should last, or click the **Data file reaches** drop down list and select the maximum file size for the recording.
6. Click **OK** to save your settings. If you selected the **When the OK button is clicked** in step 5, the recording begins.

Figure 126 *Recording Spectrum Analysis Data*



While the recording is in progress, a round, red recording icon and recording status information appears at the top of the spectrum dashboard. You can view data for other spectrum monitors and charts while the recording is in progress. If you want to stop the recording before recording period has finished, click **Stop** by the recording status information. When you the **Stop**, a popup window appears and allows you to stop and delete the current recording, stop and save the recording in its current state (before it has completed), or continue recording again.

Saving the Recording

After the recording has ended, either because the recording period has elapsed, the recording maximum file size has been reached, the **Spectrum Monitor Recording Complete** window appears and displays information for the current recording.

Figure 127 *Saving Spectrum Analysis Data*



To save the recording file:

1. From the **Spectrum Monitor Recording Complete** window, click **Continue**.
2. A **Save As** window appears and prompts you to select a file name for the recording and a location to save the file.
3. Click **Save**.

Playing a Spectrum Analysis Recording

There are two ways to play back a spectrum recording. You can use the playback feature in the spectrum dashboard, or view recordings using the Dell Networking W-RFPlayback tool downloaded from the Dell website.

Playing a Recording in the Spectrum Dashboard

The spectrum monitor does not have to be subscribed to your spectrum analysis client in order to play back a recording in the spectrum dashboard. However, you cannot play back an existing recording in the spectrum dashboard while another recording session is currently in progress.

To play a spectrum analysis recording in the spectrum dashboard:

1. Navigate to **Monitoring > Spectrum Analysis > Spectrum Dashboards** window.
2. Click the **Recording View/Play** link at the top of the window.
3. Click **Load File For Playback**.
4. An **Open** dialog box appears and prompts you to browse to and select the file you want to open.
5. Click **Open**.
6. Click the triangular play icon at the top of the window to start playing back the recording.

Recorded data for the selected spectrum monitor and dashboard view appears in the spectrum analysis dashboard. You can replace any of the graphs in the playback window with a different graph type while replaying the recording. A playback progress bar at the top of the window shows what part of the recording currently appears on the dashboard. If you pause the recording, you can click and drag the red slider on this progress bar to advance to or replay any part of the current record.

Playing a Recording Using the RFPlayback Tool

The Dell Networking W-RFPlayback tool can play spectrum recordings created in this and earlier versions of ArubaOS. Dell uses the Adobe AIR application to display spectrum recording information. If you have not done so already, follow the steps below to download and install the free Adobe AIR application and the Dell spectrum playback tool.

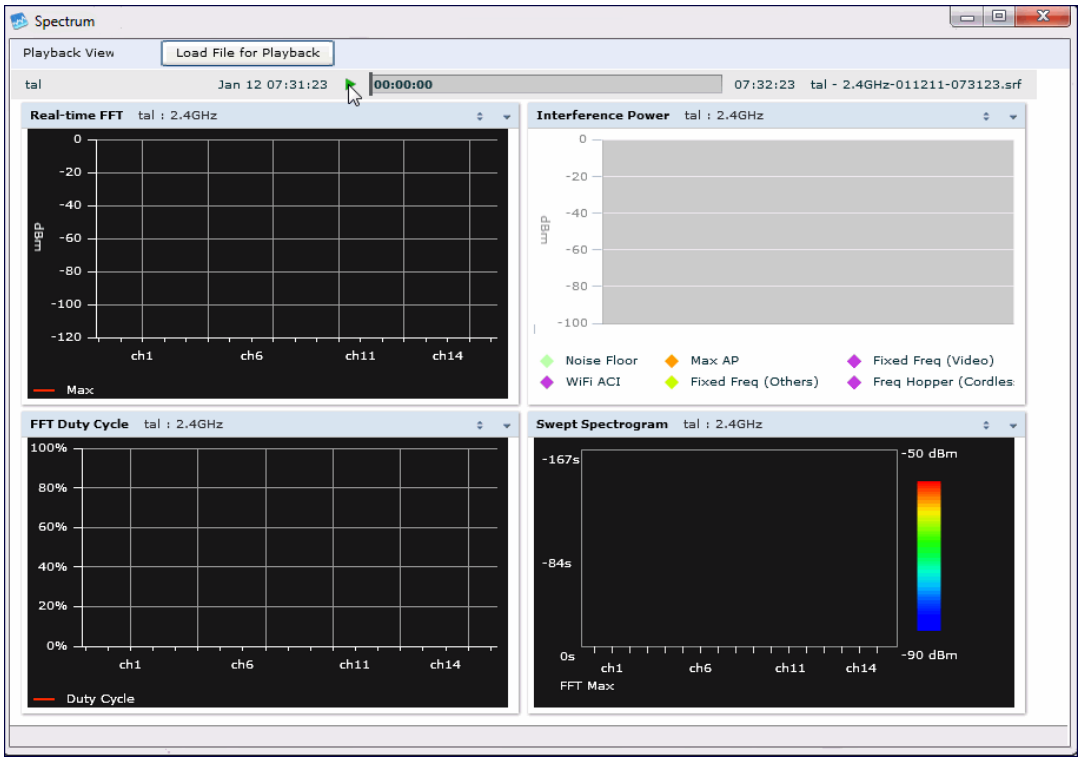
1. Download the Adobe Air application from *get.adobe.com/air/* and install it on the client on which you want to play spectrum recordings.
2. Next, download the spectrum playback installation file from the Dell website.
3. Open the folder containing the spectrum installation file, and double-click the spectrum.air icon to install the spectrum playback tool. You will be prompted to select the folder in which you want to install this tool.

Once you have installed the Dell Networking W-RFPlayback tool, follow the steps below to load and view a spectrum recording.

1. Start the Spectrum playback application.
2. Click **Load File for Playback**. An **Open** dialog box appears and prompts you to browse to and select the file you want to open.
3. Click the triangular play icon at the top of the window play the recording.

The RFPlayback tool also allows you to select and display different graph types while the recording playback is in progress. A playback progress bar at the top of the window shows what part of the recording is displayed in the playback tool. If you pause the recording, you can click and drag the red slider on this progress bar to advance to or replay any part of the current record.

Figure 128 *Playing a Recording with the Spectrum Playback Tool*



Troubleshooting Spectrum Analysis

Verifying Spectrum Monitors Support for One Client per Radio

Each spectrum monitor radio can only send information to one client at a time. If you log into a controller and the spectrum monitor dashboard does not display any data for the selected radio, another user may be logged in to the controller at that time. Note that dual-radio spectrum monitors may be accessed by two clients; one client for each radio.

Converting a Spectrum Monitor Back to an AP or Air Monitor

If want to convert a spectrum monitor radio back to AP or AM mode but the radio still comes up as a spectrum monitor, access the command-line interface and see if that spectrum monitor appears in the output of the **show ap spectrum local-override** command. If the spectrum monitor does appear in the local override profile table, issue the command **ap spectrum local-override no override ap-name <apame> spectrum-band <spectrum-band>** to remove the local override for that spectrum monitor and return the radio to AP or AM mode.

Troubleshooting Browser Issues

If you access the spectrum analysis dashboard using the Safari 5.0 browser, clicking the backspace button may return you to the previous browser screen. Avoid using the backspace button when changing dashboard view names or chart options.

If you are recording spectrum analysis data or playing back a spectrum analysis recording using a Mac client, do not minimize the browser window while the recording is in progress, as that may cause the Adobe Flash player to pause.

Loading a Spectrum View

Saved spectrum view preferences may not be backwards compatible with the spectrum analysis dashboard in earlier versions of ArubaOS. If you downgrade to an earlier version of ArubaOS and your client is unable to load a saved spectrum view in the spectrum dashboard, access the CLI in enable mode and issue the command **ap spectrum clear-webui-view-settings** to delete the saved spectrum views and display default view settings in the spectrum dashboard.

Troubleshooting Issues with Adobe Flash Player 10.1 or Later

Removing focus from the browser window displaying the spectrum analysis dashboard may cause Adobe Flash 10.1 or later to stop updating the spectrum charts to reduce CPU usage. When you restore focus to the spectrum analysis dashboard, you may see the spectrum charts update rapidly as the display catches up. Recorded data may be inaccurate if you navigate away from the spectrum window during a recording. Flash 10.0 does not have this issue.

Understanding Spectrum Analysis Syslog Messages

The spectrum analysis feature can send four different types of syslog messages: wifi add, wifi delete, non-wifi add, and non-wifi delete. All messages are in the wireless category at the syslog severity level NOTICE.

The four syslog message types appear in the following formats:

- AM: Spectrum: new wifi device found = [addr:%s] SSID = [ssid:%s] BSSID [bssid_str:%s] DEVICE ID [did:%d]
- AM: Spectrum: deleting wifi device = [addr:%s] SSID = [ssid:%s] BSSID [bssid_str:%s] DEVICE ID [did:%d]
- AM: Spectrum: new non-wifi device found = DEVICE ID [did:%u] Type [dytpe:%s] Signal [sig:%u] Freq [freq:%u]KHz Bandwidth [bw:%u]KHz
- AM: Spectrum: deleting non-wifi device = DEVICE ID [did:%d] Type [dtype:%s]

Playing a Recording in the RFPlayback Tool

The Dell Networking W-RFPlayback tool is periodically updated to support improvements to the ArubaOS Spectrum Analysis feature. The RFPlayback tool can play spectrum recordings created in the same version of ArubaOS or earlier releases. If the RFPlayback tool cannot load a newer recording, you may need to download a more recent version of the tool from the Dell website.

The ArubaOS dashboard monitoring functionality provides enhanced visibility into your wireless network performance and usage within a controller. This allows you to easily locate and diagnose WLAN issues in the controller.

The dashboard monitoring is available via the WebUI. To monitor and troubleshoot RF issues in the WLAN, click the **Dashboard** tab. The following pages in the **Dashboard** page allows you to view various performance and usage information:

- Performance
- Usage
- Security
- AppRF
- Potential Issues
- WLANs
- Access Points
- Clients
- Firewall
- AirGroup
- UCC

Additionally, you can view the context sensitive help for each field in the **Dashboard** UI by clicking the **help** link at the topmost right corner of the UI. The field for which the help has been defined appears as green. You can turn off the **help** by clicking **Done**.



You can use the **Search** functionality to find the matched results for clients, APs, and WLANs. Click the count on the search results of clients, APs, and WLANs to navigate the related summary page with the filters applied.

Performance

This page displays the performance details of the wireless clients and APs connected to the controller.

Clients

This section displays the total number of wireless clients connected to the controller. You can view the distribution of clients in different client health ranges, SNR ranges, associated data rate ranges, and data transfer speed ranges using the histograms and distributed charts. You can click on the hyperlinked number to view the data in different screens with histograms.

An AP's client health is the efficiency at which that AP transmits downstream traffic to a particular client. This value is determined by comparing the amount of time the AP spends transmitting data to a client to the amount of time that would be required under ideal conditions, that is, at the maximum Rx rate supported by client, with no data retries.

A client health metric of 100% means the actual airtime the AP spends transmitting data is equal to the ideal amount of time required to send data to the client. A client health metric of 50% means the AP is taking twice as long as is ideal, or is sending one extra transmission to that client for every packet. A metric of 25% means

the AP is taking four times longer than the ideal transmission time, or sending 3 extra transmissions to that client for every packet.

To understand histogram information, see Using Dashboard Histograms on page 748.

APs

This section displays the following performance details of the APs on the controller:

- Overall goodput
- Frame rate distribution of the APs
- Channel quality
- To client or from client frame rates
- Percentage of frames dropped

You can click the hyperlinked text and histograms to view the AP specific performance information as a trend chart. Additionally, you can view the distribution of the APs in different noise floor ranges, channel utilization ranges, and non-Wi-Fi interference ranges using the histograms. To understand histogram information, see Using Dashboard Histograms on page 748.

Using Dashboard Histograms

Dashboard histograms are a visual representation of the distribution of the wireless clients, access points, and radios across different performance parameters in the controller. Histograms help you to quickly identify any performance issues in the network from the color of the distribution. For example, critical ranges of the distribution are highlighted in red and normal ranges are highlighted in green.

You can view the number of clients or APs falling in each range of the distribution with a hyperlink. You can also perform the following tasks on the histograms to get additional information on the clients and APs in the distribution:

- **View Client or AP details**: click the hyperlinked number to view the details of the clients or APs in a pop-up window.
- **Sort**: click a column header of the clients or APs table to sort the complete list based on the entries on the active column. You can also use the sort icon that appears when you click on a column for sorting.
- **Filter**: click the filter icon and select the filter criterion on any column to filter the entries.
- **Close pop-up window**: click on the close icon to close the client or AP details pop-up window.

Usage

The Usage page displays the usage summary of the following on the controller:

- **Clients & APs**: The active wireless clients, status of APs and its usage.
- **Top APs**: The list of APs with the number of clients on the controller. The list of APs is in the descending order based on the number of clients associated with an AP. You can filter the APs for the 2.4 GHz and 5 GHz radio band options.
- **Radios**: The radios and clients connected to an AP, usage, and frame types transmitted and received by the radio.
- **Devices**: The pie chart of the clients based on the device type. Clicking on the pie chart segment opens the client details page filtered on the device type.
- **AirGroup**: All the AirGroup services available and number of servers offering the service. It is aggregated by the total number of AirGroup servers sorted by the services they advertise. For more information, see Controller Dashboard Monitoring on page 962.

- **Overall Usage**: The total number of clients and APs that have the low usage and throughput data in the last 15 minutes.
- **Usage by WLANs**: The total number of clients per WLAN and throughput data in the last 15 minutes. You can view only three WLANs in a graph and the remaining WLANs are displayed in other graph. Click the graph to view the blown up chart and information on the **Clients** page.
- **Apps by Usage:** The charts with the list of application based on the usage. You can click on the specific chart to view the application details in the **Firewall Application** page.
- **Apps by Sessions**: The list of top five applications with the session information in descending order.
- **Call Quality vs. Client Health**: This is a new graph added in ArubaOS 6.4. This graph displays the co-relation between the VoIP call quality and the VoIP client health of every Unified Communication and Collaboration (UCC) call. For more information, see UCC Dashboard in the WebUI on page 908.
- **Top Sessions**: The top five sessions by user with usage details.
- **Collaboration Apps**: The list of applications with sessions and usage details.

You can click the hyperlinked text in the sections above to view the lists and trend chart in the last 15 minutes and summary of the APs and clients in the new windows. For more information on the columns, you can view the context sensitive help for each field in the **Dashboard** UI by clicking the help link at the top right corner of the UI.

Security

This page allows you to monitor the detection and protection of wireless intrusions in your network.

The two top tables—**Discovered APs & Clients** and **Events**—contain data as links. When these links are selected, they arrange, filter, and display the appropriate information in the lower table.



The term **events** in this document refers to security threats, vulnerabilities, attacks (intrusion or Denial of Service), and other related events.

Potential Issues

This page displays the total number of radios and wireless clients that may have potential issues in the network. You can click on the total number to view the trend of the clients and radios with potential issues in the last 15 minutes. You can also view the number of clients or radios that have a specific potential issue in each radio band.

The potential issues that a client may have are:

- **Low SNR**: clients that have signal to noise ratio of 30 dBm or lower.
- **Low speed**: clients that have a connection speed of 36 Mbps or lower.
- **Low goodput**: clients that have an average data rate of 24 Mbps or lower.

The potential issues that a radio may have are:

- **High noise floor**: radios that have a noise floor of -85 dB or greater.
- **Busy channel**: radios that have a channel utilization of 80% or greater.
- **High non-Wi-Fi interference**: radios that have a non-Wi-Fi interference of 20% or greater.
- **Low goodput**: radios that have an average data rate of 24 Mbps or lower.
- **High client association**: radios that have 15 or more clients connected.

You can click on the hyperlinked number to view the details of the respective clients or radios in the bottom pane of the page. You can perform the following tasks on the details table:

- **Sort**: click a column header of the table to sort the complete list based on the entries on the active column. You can also use the sort icon that appears when you click on a column for sorting.
- **View or hide columns**: click the drop-down menu on the top right corner of the table header and select **Custom Columns**; choose the **Edit Current View** option to select the columns that you want to view.

WLANs

You can view the WLAN details such as the number of associated APs, radios, wireless clients, and the WLAN usage in the controller. You can also view the details of the associated APs and clients as tables.

The following sections are available in the WLANs page:

- **WLANs**: the unique SSID of the WLAN, clients connected in the network, APs connected to the WLAN, Radios that are enabled on the AP, Goodput, usage, and the frames transmitted and received by the AP.
- **All WLANs**: the clients, usage, and device distribution information in graphs.

Click the hyperlinked text in the WLANs page to view the following menus with the summary:

- **Info**: the summary of the WLAN details, frames transmitted and received from and to the client, air quality, and Tx/Rx statistics.
- **Clients**:the summary of WLANs and clients.
- **Radios**: the summary of APs and clients, channel, and its utilization.
- **Charts**: the summary of WLAN details in graphs.
- **Firewall**: the summary of users, destination, applications, devices and its roles.

You can perform the following tasks on this page:

- **Sort**: click a column header of the WLAN table to sort the complete list based on the entries on the active column. You can also use the sort icon that appears when you click on a column for sorting.
- **Filter**: click the filter icon and select the filter criterion on any column of the details table to filter the entries.
- **Customize column view**: click the drop-down menu on the top right corner of the table header and select **Custom Columns**; choose the **Edit Current View** option to select the columns that you want to view. You can also choose one of the following system defined views that have the appropriate pre-selected columns.
  - Default Columns: you cannot edit this view.
  - To/From Client Stats: you can customize this view using the **Edit Current View** option.
- **View WLAN trends**: the trends of the clients connected in the WLAN and the WLAN usage in the last 15 minutes.
- **View client summary**: click on the hyperlinked client name on the client details table to view the **Client Summary** page. In this page, you can view the client details summary (air quality metrics and from and to clients statistics), bandwidth of the client usage, trend of the client frame loss in the last 15 minutes, and the frame rate distribution of the client.
- **View AP or radio summary**: click on the hyperlinked AP name or the radio band on the AP details table to view the **Access Points** page. In this page you can view the summary of the AP details such as air quality metrics, from and to clients statistics, and the number of clients associated with the AP under different SNR ranges. Additionally, you can view the details of the associated clients and WLANs.

Access Points

You can view the details of all the radios and APs associated with the controller by selecting the specific section. You can also view the trends of the connected wireless clients and the client usage under the 2.4 Ghz and 5 Ghz radio bands in the last 15 minutes.

The **Access Points** page has the following three sections:

- **Access Points**—Displays the AP name, status, uptime, mode, and model details.
- **Radios**—Displays the AP name, band, radio mode, goodput, usage, and the frames transmitted and received by the AP.
- **All Clients**—Displays the clients and usage trend in charts for the last 15 minutes.

You can click the hyperlinked text on the **Access Points** page to view the following menus with the summary:

- **Info**—Displays the summary of the AP details, frames transmitted and received from and to the client, air quality, and Tx/Rx statistics.
- **WLANs & Clients**—Displays the summary of WLANs and clients.
- **Charts**—Displays the summary of clients and its usage in graphs for different bands.
- **History**—Displays the history of channel utilization, frame drops, and frame rates for every minute with histograms for the last 15 minutes.

You can perform the following tasks on this page:

- **Sort**: Click a column header of the AP table to sort the complete list based on the entries on the active column. You can also use the sort icon that appears when you click on a column for sorting.
- **Filter**: Click the filter icon and select the filter criterion on any column of the details table to filter the entries.
- **Customize column view**: Click the drop-down menu on the top right corner of the table header and select **Custom Columns**; choose the **Edit Current View** option to select the columns that you want to view. You can also choose one of the following system defined views that have the appropriate pre-selected columns.
  - Default Columns—You cannot edit this view.
  - Air Quality Metrics—You can customize this view using the **Edit Current View** option.
  - To/From Client Stats—You can customize this view using the **Edit Current View** option.
- **View client details**: Click on the number of clients associated with the AP to view the details of the clients on the **Clients** page.
- **View AP or radio summary**: Click on the hyperlinked AP name or the radio band on the AP details table to view the summary of the AP details such as air quality metrics, from and to clients statistics, and the number of clients associated with the AP under different SNR ranges. Additionally, you can view the details of the associated clients and WLANs.

Clients

You can view the details of all the wireless clients on the controller. You can also view the trends of the connected clients and the client usage under the 2.4 Ghz and 5 Ghz radio bands in the last 15 minutes.

The Clients page displays the following sections:

- **Clients:** The connectivity type, radios, client health, goodput, channel, and the frames transmitted and received.
- **All Clients**: The clients and its usage for 2.4 GHz and 5 GHz bands.

Click the hyperlinked text on the Clients page to view the following menus with the summary:

- **Info**: The summary of the client details, frames transmitted and received from and to the client, air quality, and Tx/Rx statistics.
- **Charts**: The summary of the client details in graphs.
- **AirGroup**: A list of all the far and near end devices that are either accessible or not accessible by the specific client.For more information, see Controller Dashboard Monitoring on page 962.
- **Firewall**: The summary of traffic in the clients, applications and its roles, and protocols.

- **UCC**: This tab displays an aggregated list of UCC call data metrics of a client. For more information, see UCC Dashboard in the WebUI on page 908.



The AirGroup and Firewall links are not available on W-600 Series controllers.

You can perform the following tasks on this page:

- **Sort**: Click a column header of the AP table to sort the complete list based on the entries on the active column. You can also use the sort icon that appears when you click on a column for sorting.
- **Filter**: Click the filter icon and select the filter criterion on any column of the details table to filter the entries.
- **Customize column view**: Click the drop-down menu on the top right corner of the table header and select **Custom Columns**; choose the **Edit Current View** option to select the columns that you want to view. You can also choose one of the following system defined views that have the appropriate pre-selected columns.
  - Default Columns: you cannot edit this view.
  - Air Quality Metrics: you can customize this view using the **Edit Current View** option.
  - To/From Client Stats: you can customize this view using the **Edit Current View** option.
- **View client summary**: Click on the hyperlinked client name on the client details table to view the **Client Summary** page. In this page, you can view the client details summary (air quality metrics and from or to clients statistics), bandwidth of the client usage, trend of the client frame loss in the last 15 minutes, and the frame rate distribution of the client.
- **View AP details**: Click on the hyperlinked AP name to view the **Access Points** page.
- **View WLAN details**: Click on the hyperlinked SSID of the WLAN to view the **WLANs** page.

Firewall

The ArubaOS Policy Enforcement Firewall (PEF) module provides identity-based controls to enforce application-layer security and prioritization. With PEF, network administrators can enforce network access policies that specify who may access the network, with which mobile devices, and which areas of the network they may access. The Dell AppRF technology integrated with PEF delivers mobile application traffic visibility through a simple dashboard that shows the applications in use by user and device. It gives network administrators insights on the applications that are running on their network, and the users using them.

The **Firewall** page on the **Dashboard** tab displays the PEF summary of all the sessions in the controller aggregated by users, devices, destinations, applications, WLANs, and roles.

Firewall visibility is disabled on the controller by default. To enable this feature, use the following procedures:

In the WebUI

1. Navigate to the **Dashboard > Firewall** page.
2. Click the link on the **Element View** section to enable firewall visibility. To disable, click the **Disable Firewall** link at the bottom of the **Element View** section.

In the CLI

Use the following command:

```
(host)(config) #firewall-visibility
```

To disable this setting, include the `no` parameter:

```
no firewall-visibility
```

This feature is supported in W-3000 Series, W-6000 controllers, and requires the PEFNG license. For W-7200 Series controller, see AppRF on page 758.

Element View

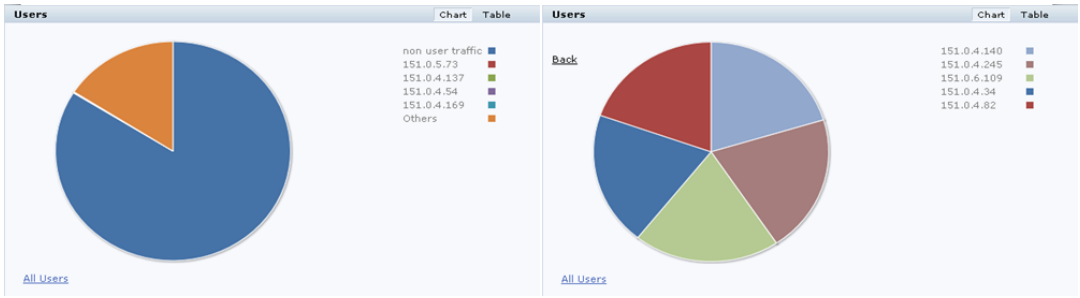
Navigate to the **Dashboard > Firewall** page to view **Element View** section. This section displays a summary of all the sessions in the controller and includes six categories of monitoring data, or elements, that display traffic statistics aggregated by the following elements:

Table 151: *Element View*

Element	Description
User	Indicates a wireless or wired user associated to the controller. Traffic that is not generated by a user is aggregated as **non-user traffic**.
Devices	Specifies the client device type. for example: Windows 7, Mac OS X, iPhone, or Android.
Destinations	Destination hostname, or IP address if the hostname is unavailable. Common advertising and file sharing services on the Internet are categorized under special destinations called **ad networks** and **file share networks** respectively.
Applications	Application name, protocols, and ports. For example: • **Web applications**: YouTube, Twitter, Facebook, Gotomeeting, Webex, Amazon, Saleforce, and more. • **Stateful applications**: FTP, Lync, SIP, and more. • **Custom applications**: using the `netservice` command, you can define custom applications if the application uses well-known port numbers (0 - 1023). • **Peer-to-Peer**: all peer-to-peer traffic is classified under **peer to peer**. • **Lync applications**: Lync-desktop-sharing, Lync-file-transfer, Lync-voice, and Lync-video. If a session does not map to any of the above, the destination port is classified as **application**.
WLANs	The service set identifier (SSID) that uniquely identifies the WLAN. Wired connection is shown as **wired**.
Roles	Determines the user's network privileges based on the assigned user role.

The **Element View** section has two views: **Chart** and **Table**. Click **Chart** or **Table** at the top-right corner of an element to toggle between the two views. Each chart container shows the top five sessions with respect to traffic bandwidth and the rest are shown as **Others**. Click **Others** within the chart to view the rest of the sessions in the chart. Click any entry on the chart legend to view more usage details. The figure below shows the **Chart** view:

Figure 129 *Chart View*



In addition to the element, the **Table** view shows the common fields displayed in the table below:

Table 152: *Table View Fields*

Column	Description
Bytes	Total number of bytes transmitted and received by an element.
Tx Bytes	Total number of bytes transmitted by an element.
Rx Bytes	Total number of bytes received by an element.

You can perform the following tasks in the **Table** view:

- **Sort**: click a column header of the table to sort the list by column. You can also use the sort icon that appears when you click on a column.
- **Filter**: click the filter icon on the first column and select the filter criterion to filter the entries.

Details View

Navigate to the **Dashboard > Firewall** page. Click the **All <element>** link to view the **Details View** page. There are four sections on this page.

Element Tab

The **Element Tab** shows the available usage detail elements. Click an element to view more usage details:

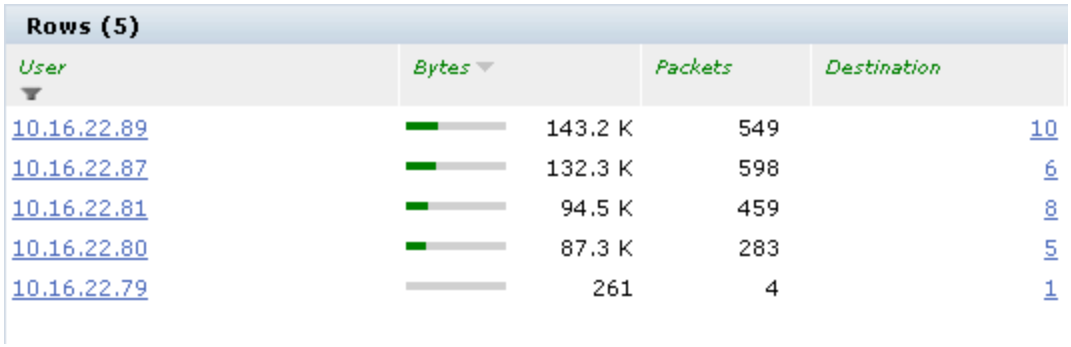
Figure 130 *Element Tab*



Element Summary View

The **Element Summary View** displays a detailed view of all the six elements and their corresponding fields:

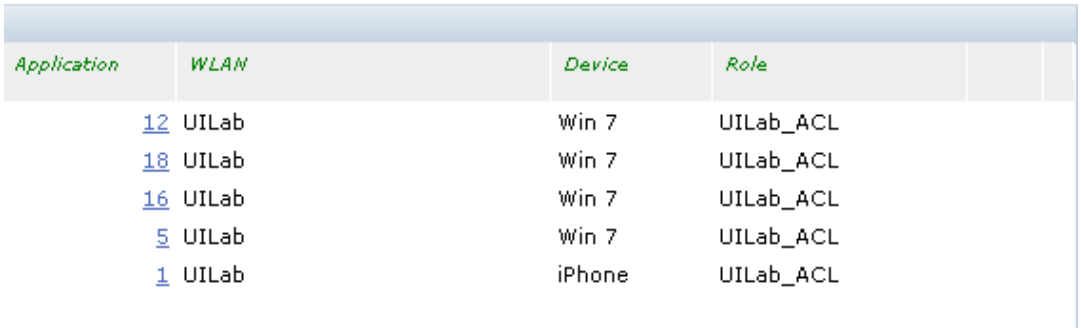
Figure 1a *Element Summary View*



Rows (5)

User	Bytes	Packets	Destination
10.16.22.89	143.2 K	549	10
10.16.22.87	132.3 K	598	6
10.16.22.81	94.5 K	459	8
10.16.22.80	87.3 K	283	5
10.16.22.79	261	4	1

Figure 1b *Element Summary View (continued)*



Application	WLAN	Device	Role
12	UILab	Win 7	UILab_ACL
18	UILab	Win 7	UILab_ACL
16	UILab	Win 7	UILab_ACL
5	UILab	Win 7	UILab_ACL
1	UILab	iPhone	UILab_ACL

See the following table for more information on **Element Summary View** fields:

Table 153: *Element Summary View Fields*

Column	Description
User	Indicates a wireless or wired user associated to the controller. Click a **User** IP address to view details of the connected client.
Bytes	Total number of bytes transmitted and received by an element.
Packets	Total number of data packets transmitted and received by an element.
Device	Specifies the client device type. Click the number to view details of the device type identification.
Destination	Total number of destination hostnames or IP addresses. Click the number to view details of the destination hosts.
Application	Total number of application name, protocols, and ports. Click the number to view details of the application ports.
WLAN	The service set identifier (SSID) that uniquely identifies the WLAN. Click the number to view details of the WLAN SSID.
Role	Determines the user's network privileges based on the assigned user role. Click the number to view details of the role.

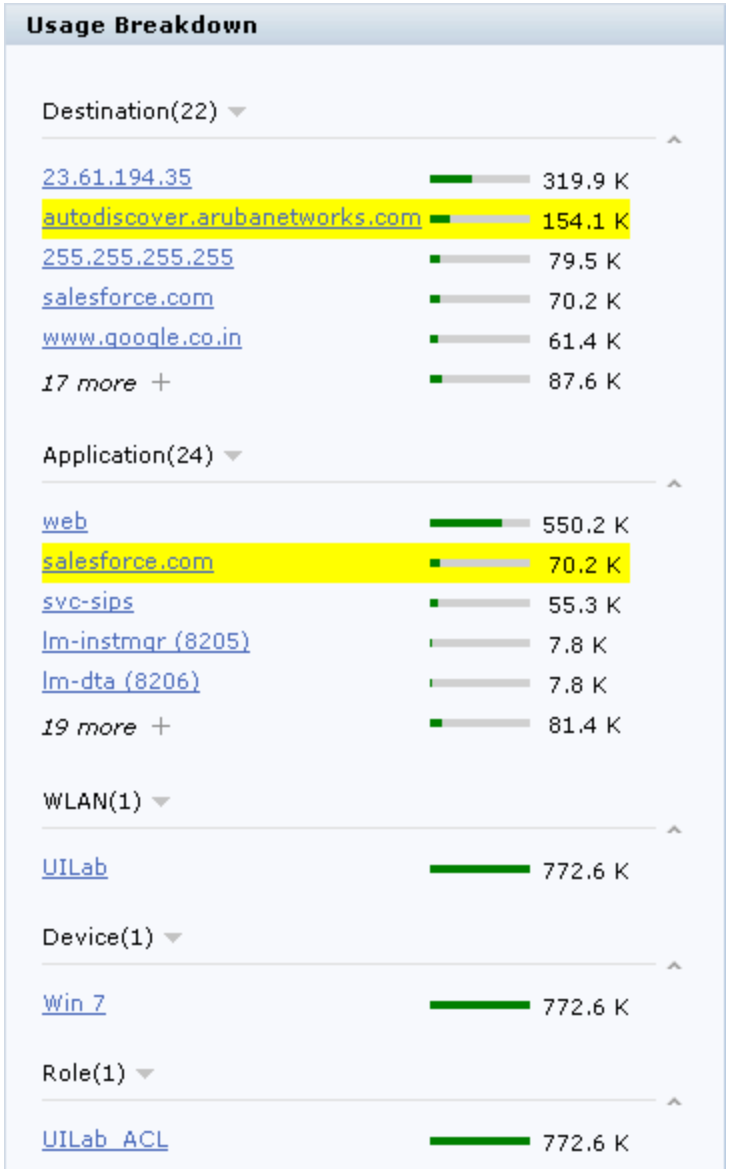
You can perform the following tasks in the **Element Summary View**:

- **Sort**: click a column header of the table to sort the list by column. You can also use the sort icon that appears when you click on a column.
- **Filter**: click the filter icon on the first column and select the filter criterion to filter the entries.

Usage Breakdown

In the **Usage Breakdown,** section you can apply any of the filters that are listed under each element to customize the output. To apply a filter, click any row under each element. The selected row turns yellow. The filtered output is displayed in the **Element Summary View** and **Aggregated Sessions** sections of the page. Click the row again to deselect it and remove the filter. For example, if you click **autodiscover.arubanetworks.com** under **Destination**, and **salesforce.com** under **Application**, the **Element Summary View** and **Aggregated Sessions** sections display session information based on the selected rows. The following figure shows the selected row in each element:

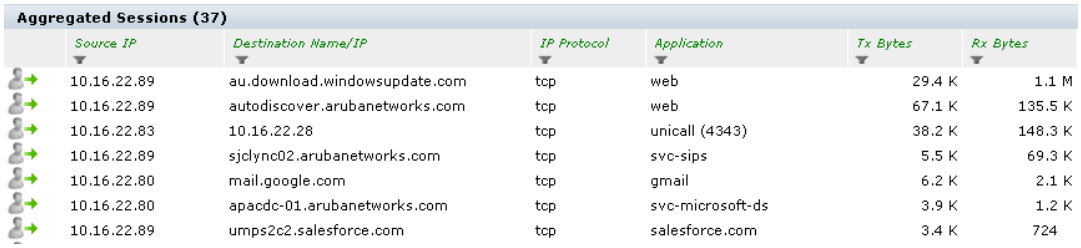
Figure 131 *Usage Breakdown*



Aggregated Sessions

The **Aggregated Sessions** displays a list of all user and non-user sessions on the controller.

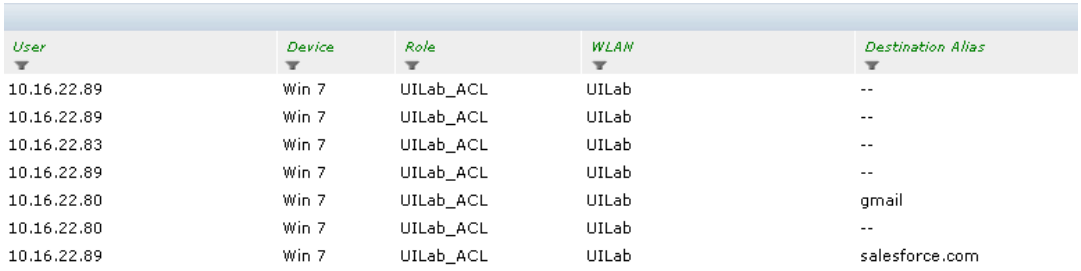
Figure 2a *Aggregated Sessions*



Aggregated Sessions (37)

Source IP	Destination Name/IP	IP Protocol	Application	Tx Bytes	Rx Bytes
10.16.22.89	au.download.windowsupdate.com	tcp	web	29.4 K	1.1 M
10.16.22.89	autodiscover.arubanetworks.com	tcp	web	67.1 K	135.5 K
10.16.22.83	10.16.22.28	tcp	unicall (4343)	38.2 K	148.3 K
10.16.22.89	sjclync02.arubanetworks.com	tcp	svc-sips	5.5 K	69.3 K
10.16.22.80	mail.google.com	tcp	gmail	6.2 K	2.1 K
10.16.22.80	apacdc-01.arubanetworks.com	tcp	svc-microsoft-ds	3.9 K	1.2 K
10.16.22.89	umps2c2.salesforce.com	tcp	salesforce.com	3.4 K	724

Figure 2b *Aggregated Sessions (continued)*



User	Device	Role	WLAN	Destination Alias
10.16.22.89	Win 7	UILab_ACL	UILab	--
10.16.22.89	Win 7	UILab_ACL	UILab	--
10.16.22.83	Win 7	UILab_ACL	UILab	--
10.16.22.89	Win 7	UILab_ACL	UILab	--
10.16.22.80	Win 7	UILab_ACL	UILab	gmail
10.16.22.80	Win 7	UILab_ACL	UILab	--
10.16.22.89	Win 7	UILab_ACL	UILab	salesforce.com

See the following table for more information on **Aggregated Sessions** fields:

Table 154: *Aggregated Sessions Fields*

Column	Description
Source IP	Indicates the IP address of the wireless or wired user associated to the controller.
Destination Name/IP	Destination hostname, or IP address if the hostname is unavailable.
IP Protocol	Type of IP protocol traffic: for example, TCP or UDP.
Application	Application name, protocols, and ports.
Tx Bytes	Total number of bytes transmitted in a session.
RX Bytes	Total number of bytes received in a session.
User	Indicates a wireless or wired user associated to the controller.
Device	Specifies the client device type.
Role	Determines the user's network privileges based on the assigned user role.
WLAN	The service set identifier (SSID) that uniquely identifies the WLAN.
Destination Alias	Fully Qualified Domain Name (FQDN) or the URL of the destination network or host.

You can perform the following tasks in the **Aggregated Sessions** section:

- **Sort**: click a column header of the table to sort the list by column. You can also use the sort icon that appears when you click on a column.
- **Filter**: click the filter icon on the first column and select the filter criterion to filter the entries.

AppRF

AppRF is a application visibility and control feature and was introduced in ArubaOS 6.4.0, AppRF does deep packet inspection (DPI) of local traffic, and detects over 1500 applications on the network. AppRF allows you to configure both application and application category policies within a given user role.

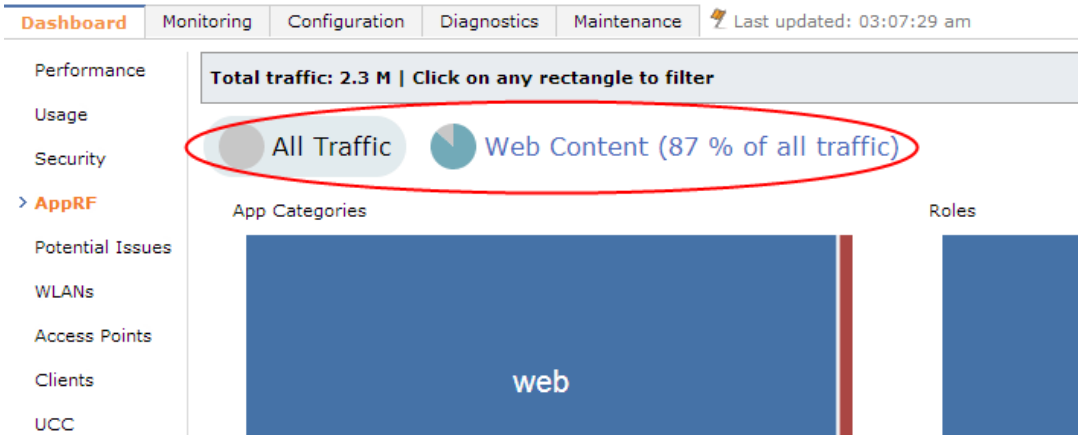
NOTE

The AppRF dashboard application visibility feature is supported in only W-7200 Series controllers and W-7000 Series controllers, and requires the PEF-NG license.

In the WebUI, the AppRF dashboard contains the following two pages as shown in Figure 132:

- All Traffic: The All Traffic page displays the summary of all traffic in the controller. This is the default page. For more details, see All Traffic on page 758.
- Web Content: The Web Content page link include the percentage of all traffic in parenthesis. The Web Content page displays the summary of only the web traffic in the controller. For more details, see Web Content Classification on page 766.

Figure 132 *All Traffic and Web Content Page Options*



All Traffic

The **All Traffic** page on the **Dashboard** > **AppRF** page displays the PEF summary of all the sessions in the controller aggregated by users, devices, destinations, applications, WLANs, and roles.The applications, application categories, and other containers are represented in box charts instead of pie charts.

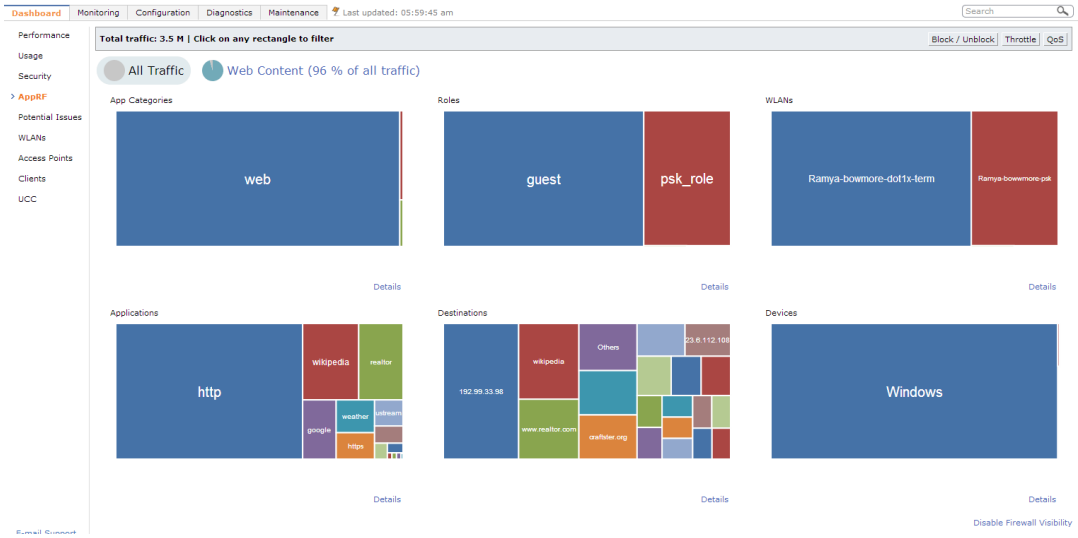
Enable DPI to enhance the benefit of the existing visualization or dashboard, To enable DPI, see the Enabling Deep Packet Inspection (DPI) section.

To view the AppRF dashboard, in the WebUI:

1. Navigate to the **Dashboard > AppRF** page.
2. Click on the link **To enable this feature, click here** to enable firewall visibility. To disable, click the **Disable Firewall Visibility** link at the bottom the page.

You will see a screen similar to the following figure.

Figure 133 *AppRF- All Traffic Page*



Action Bar

The Action bar displays the total traffic depending on the filters applied, allows the user to configure per Application, per Role, and Global Policy, and includes Action buttons namely, Block/Unblock, Throttle, and QoS.

Figure 134 *Action Bar*

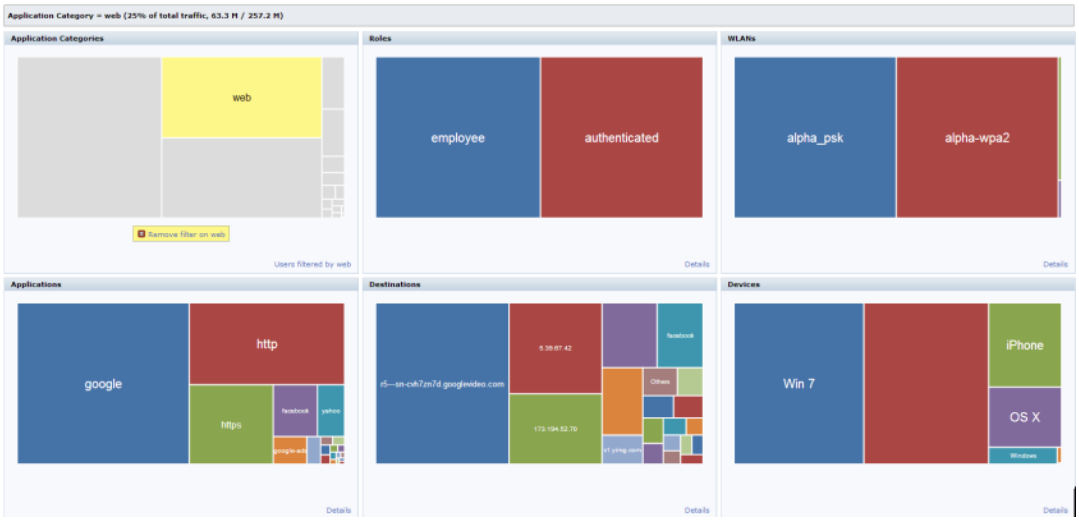


Filters

You can click on any rectangle tile in a container and that filter is applied across all the containers.

For example: If you click on the **Web** rectangle in the **Application Categories** container, Application Categories == Web filter is applied to all other containers (Roles, WLANs, Application, Destination and Devices). See the following figure.

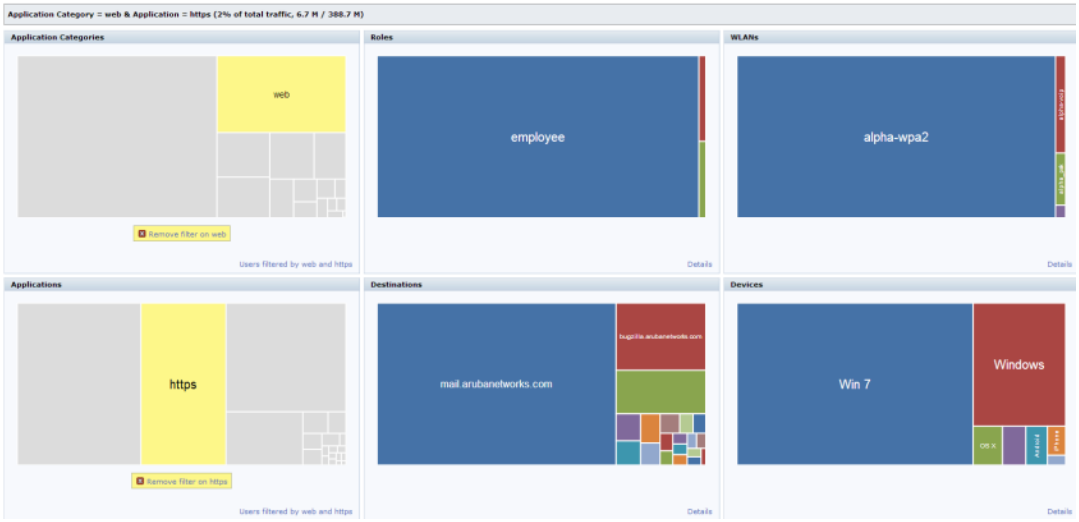
Figure 135 *Single Filter Applied*



You can apply multiple filters from different containers by clicking on muliple rectangle tiles in various containers.

For example: If you click on the **Web** rectangle in the **Application Categories** container and the **https** rectangle under **Application**, the remaining containers (Roles, WLANs, Destination and Devices) will be filtered on Application categories == web and Application == https. See the following figure.

Figure 136 *Multiple Filters Applied*



The action bar reflects the total traffic based on the filter applied. For example, see Figure 137 and Figure 138.

Figure 137 *Total traffic with Web Filter*

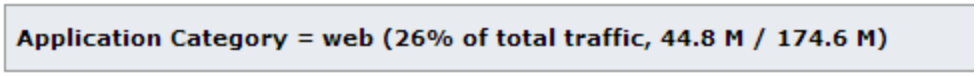
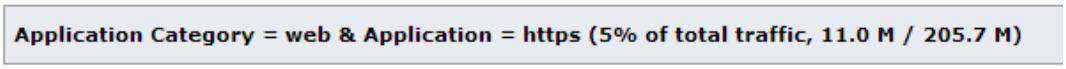


Figure 138 *Total traffic with Web and https Filter*



NOTE

The action buttons are disabled if the applied filter contains anything apart from Role and Application or Role and Application category.

To remove filters, click on **Remove filter** in the container that filter is removed across all the containers.

Details

Clicking on **Details** navigates you to the corresponding details page with data filtered by all selected rectangle when a filter is applied, The **Details** link changes to **User filtered by <filter>** in that container. See Figure 139 and Figure 140

Figure 139 *Details*

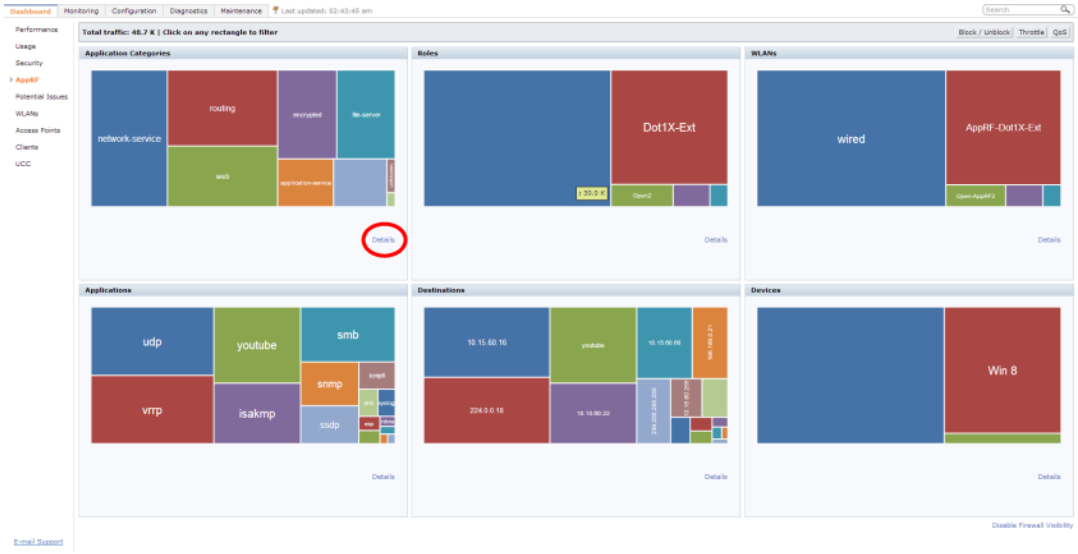
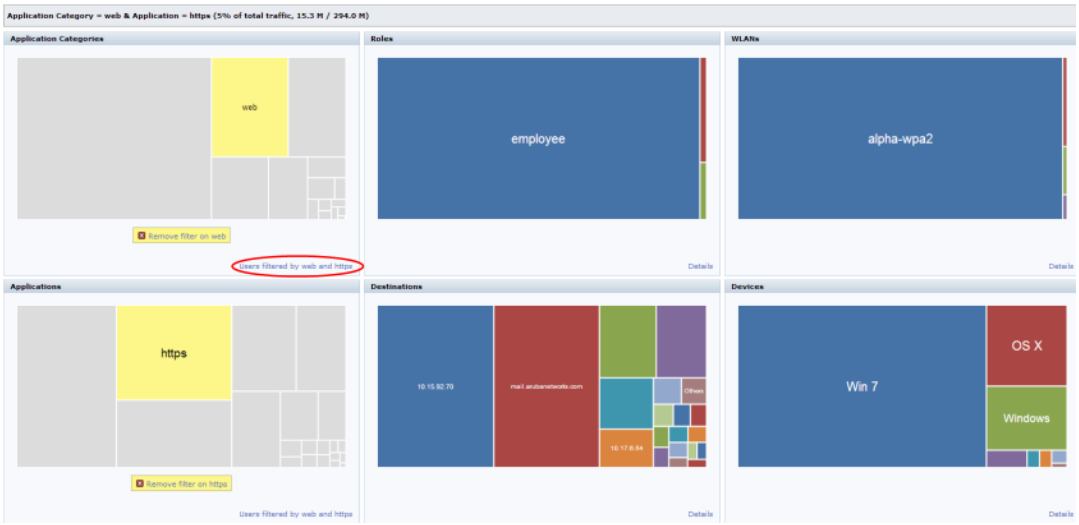


Figure 140 *User filtered by <filter>*



Clicking on **Details** or **User filtered by <filter>** shows the user table, See Figure 141 and Figure 142.

Figure 141 *Details View*

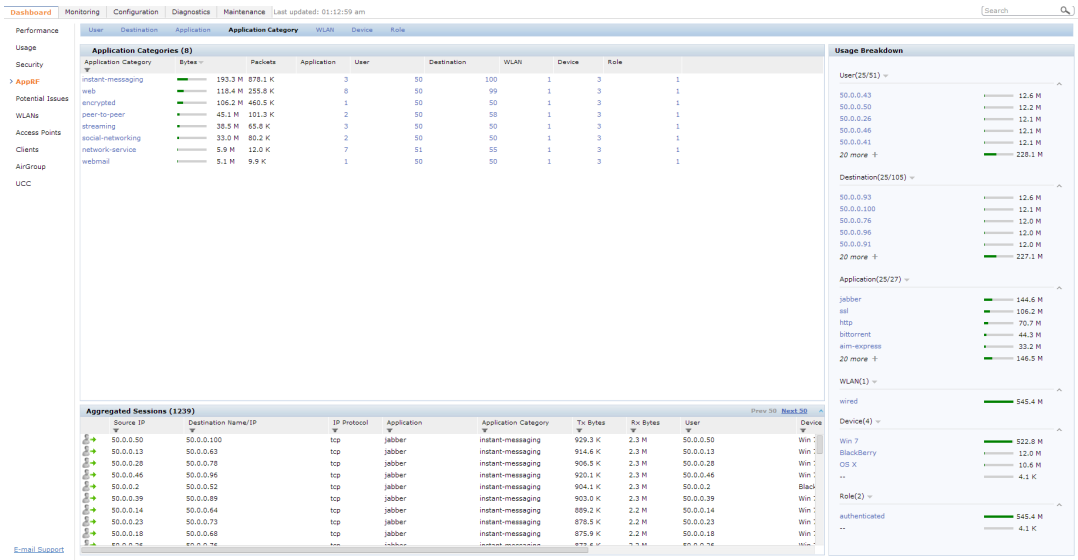
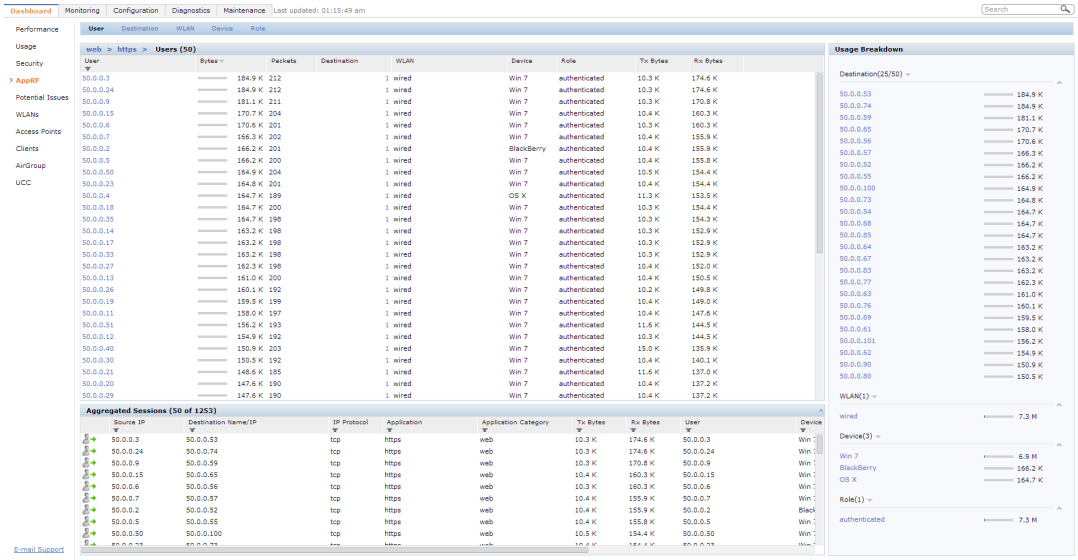


Figure 142 *User filtered by <filter>- Details View*



Block/Unblock, Throttle, and QoS Action Buttons

The pop-up window that is displayed for block/unblock, throttle, or QoS depends on the filters applied..

Upon clicking **OK**, the corresponding CLI commands are executed and the pop-up window closes retaining the filters in the AppRF main page. When filters are not applied, all the pop-up windows allow the user to configure global or per –role configuration.

The following table shows the pop-up window with respect to the Action button and the filter applied:

Table 155: *Pop-up Window with Respect to the Action Button and the Filter Applied*

Action Button	Filter	Config Level
Block/Throttle/QoS	Non-application/role ex: WLANS	No pop-up
Block	No Filters	Global and per role
Block	Application	Global
Block	Application Category	Global
Block	Application and Role	Global and per role
Block	Application category and Role	Global and per role
Throttle	No	Global and per role
Throttle	Application	Global
Throttle	Application Category	Global
Throttle	Application and Role	Global and per role
Throttle	Application category and Role	Global and per role
QoS	Application	Global
QoS	Application Category	Global

Block/Unblock

This button allows you to permit/deny an application or an application category for a given role. You can create global and per-role rules.

For example, you can block the YouTube application,which belongs to the Streaming application category for the guest role within the enterprise.

Applying a New Rule Using AppRF

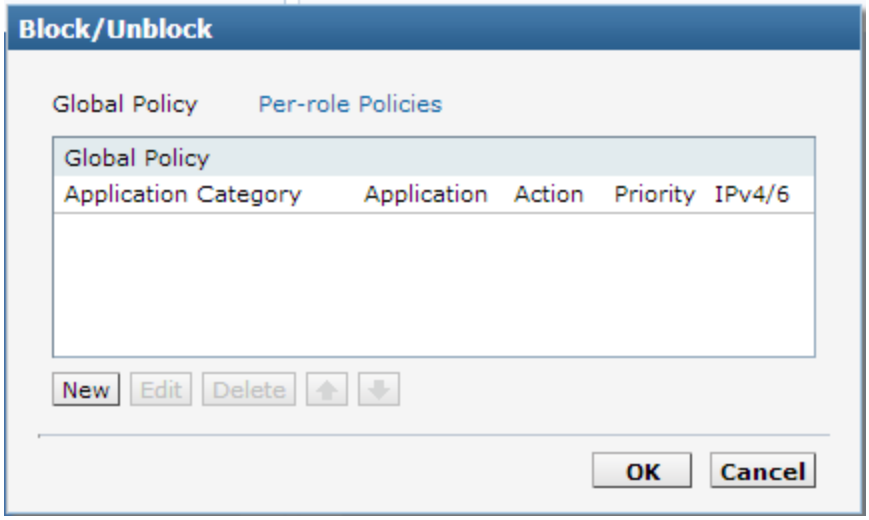
1. Click on **Block/Unblock** on the Action bar.



The **Block/Unblock** button changes to the **Block** button if a filter is applied. The pop-up window appears based on the filters applied is shown in Table 155. Click on **Show policy tables**. **Block** allows only permit action and priority setting.

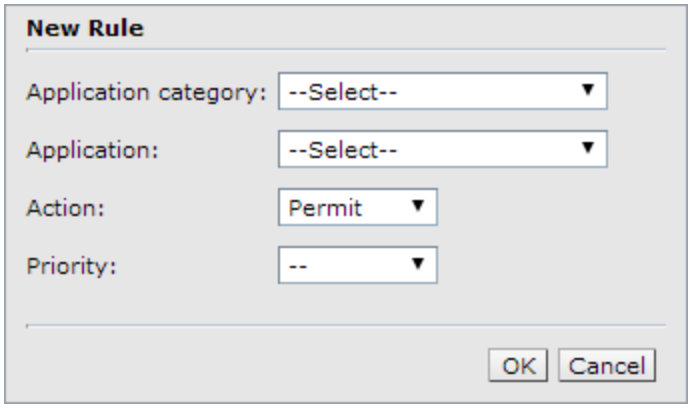
2. To create a new Global rule:
   a. Click on the **Global Policy** tab, the following pop-up window appears:

Figure 143 *Global Policy Tab*



b. Click on **New.** The following pop-up window appears:

Figure 144 *New Rule Pop-up Window*

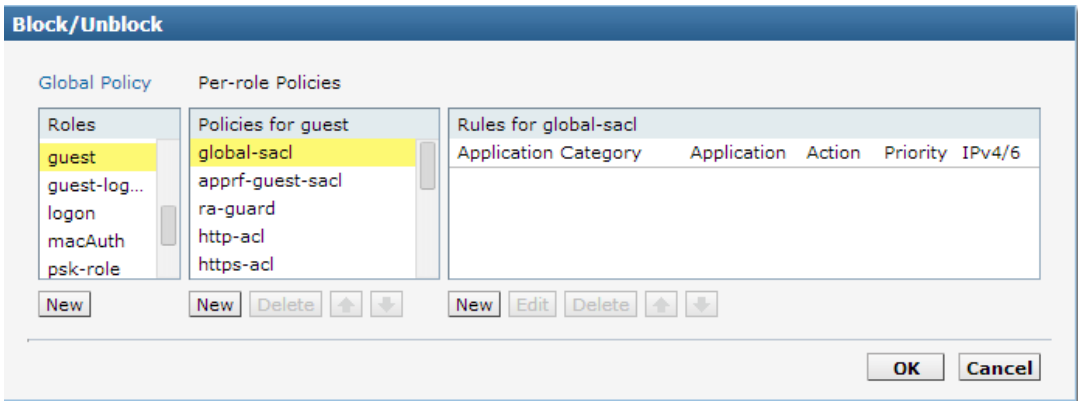


c. Select an **Application category**, **Application**, **Action**, and **Priority**.

3. To create a new per-role rule:

a. Click on the **Per-role policies** tab, the following pop-up window appears:

Figure 145 *Per-role Policies Tab*



b. Select a role from the list, or click on **New** below the role pane to create a new role and select the newly created role.

c. Select a policy from the list, or click on **New** below they policy pane to create a new policy and select the newly created policy.

d. Select an **Application category**, **Application**, **Action**, and **Priority from the New Rule pop-up window, as shown in** Figure 144

4. Click on **OK.**

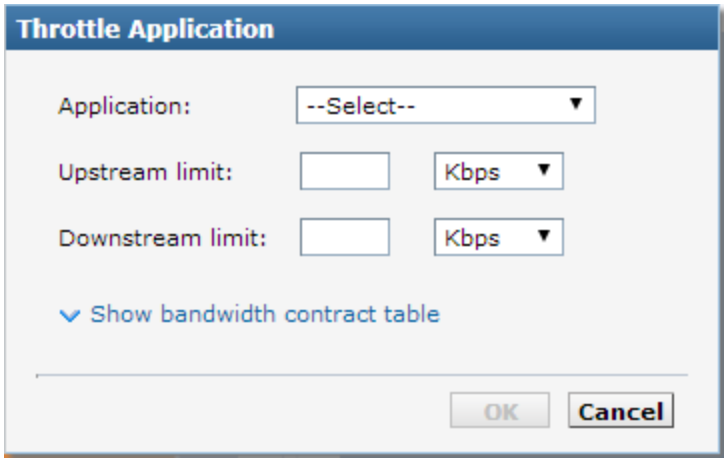
Throttle

This button allows you to limit the bandwidth usage of an application or an application category on a given role. So, you can set the upstream limit and downstream limit for an application or an application category on a given role.

For example, you can rate limit applications video streaming applications like YouTube, Netflix.

You can also view the bandwidth contract table and create a new bandwidth contract.See the following figure.

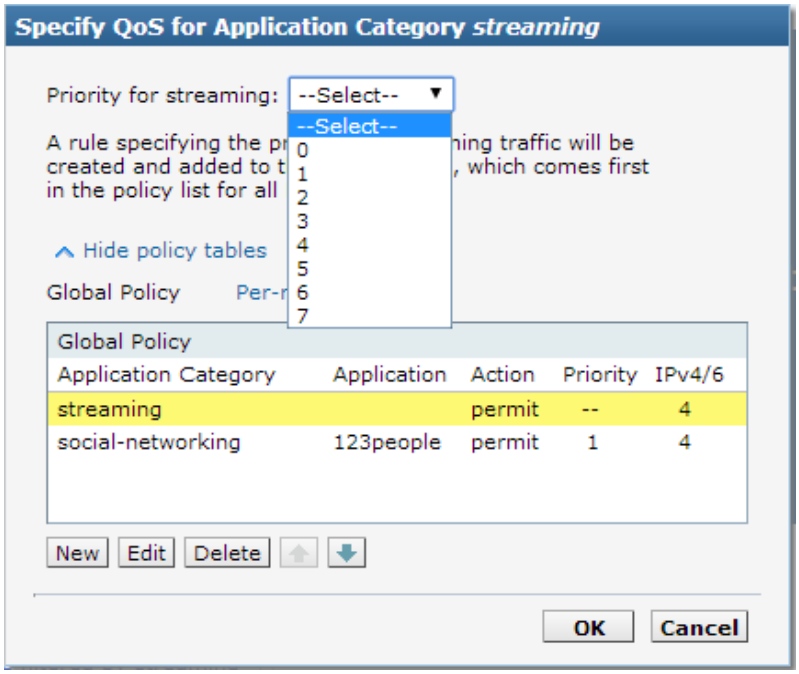
Figure 146 *Throttle Application and New Bandwidth Contract*



QoS

This button allows you to set the priority for a given application or an application category on a given role. For example, you can set the video/voice sessions originating from wireless users with a different priority to that of other web applications so that traffic would be prioritized accordingly in the your network.

Figure 147 *QoS for Application Category Streaming*



Web Content Classification

Many applications are moving to the web and web being so dynamic in nature, ArubaOS 6.4.2.0 introduces web content control through the Web Content Classification (WebCC) feature. WebCC uses a cloud-based service to dynamically determine the types of websites being visited, and their safety.

NOTE

This feature is available for all customers with a PEF license to use during an early preview period. Eventually, Dell intends to license this feature as an annual subscription. License enforcement timeline and pricing information will be made available once the SKUs and prices are finalized.

The implementation of WebCC feature can be viewed on this new web page.

When the WebCC feature is enabled, all web traffic (http and https) is classified. The classification is done in data path as the traffic flows through the controller and updates dynamically.

Dell has partnered with Webroot®, and uses the Webroot's URL database and the cloud look-up service to classify the web traffic. Dell uses Webroot classified categories and score for web categories and reputation for WebCC.

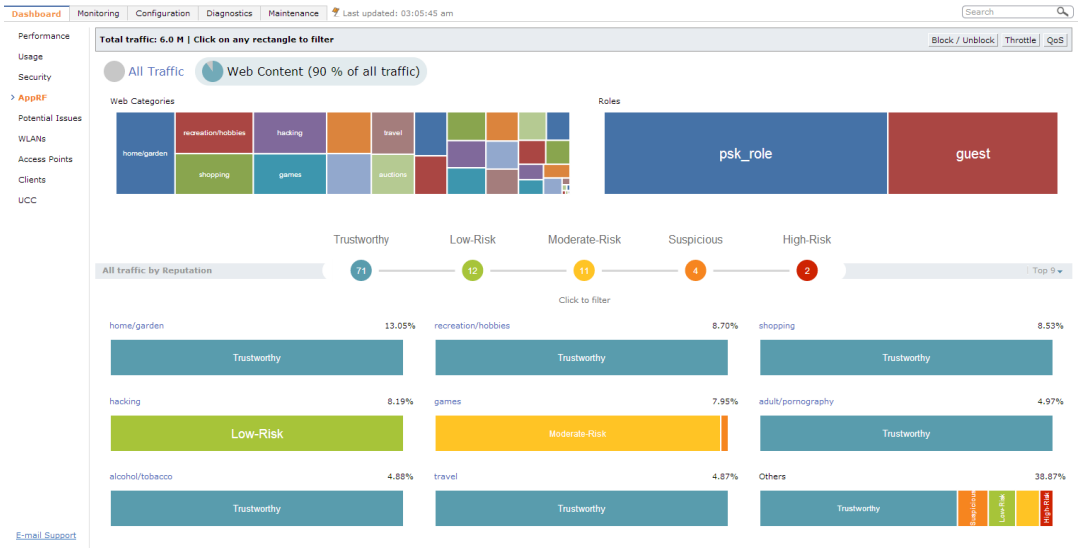
The current policy enforcement model in Dell relies on L3/L4 information of the packet or L7 information with Deep Packet Inspection (DPI) support to apply rules. WebCC complements this as the user is allowed to apply firewall policies based on web content category and reputation.

Benefits of WebCC:

1. Prevention of malicious malware, spyware, or adware by blocking known dangerous websites
2. Visibility into web content category-level
3. Visibility into web sites accessed by the user

To view the web content page from the WebUI, navigate to **Dashboard** > **AppRF** . Click on **Web Content** tab. The following figure shows the Web Content page.

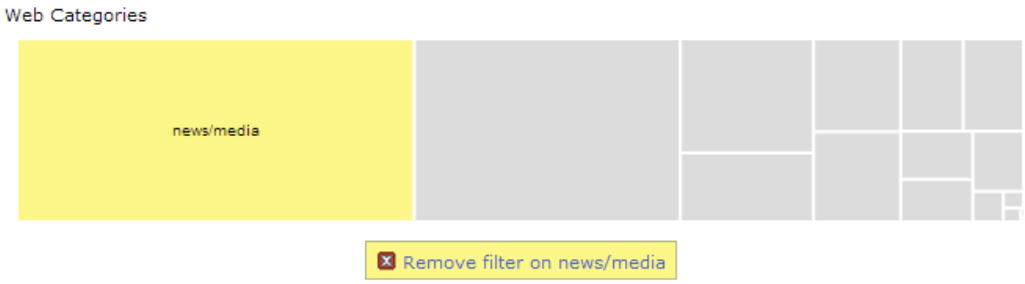
Figure 148 *Web Contents Page*



The web content page includes the following containers:

- **Web Categories:** This chart shows traffic for web categories in tree chart presentation. All boxes in this chart is click-able. Clicking on a box filters rest of page data with the clicked web category as filter, and this chart is locked until the filter is removed by clicking on **Remove filter on <web category>**. For example, see the following figure.

Figure 149 *Filter by Web Category*



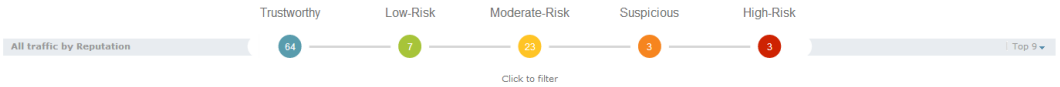
- **Roles:** This chart shows the for Roles using the web traffic in tree chart presentation. All role boxes are In this chart is click-able. Clicking on box filters rest of page data with the clicked Role as filter, and this chart is locked until the filter is removed by clicking on **Remove filter on <role name>**. For example, see the following figure.

Figure 150 *Filter by Role*



- **All traffic by Reputation:** The reputation traffic light chart shows the percentage of traffic based on reputation or score of web traffic in the controller. The reputation levels are Trustworthy, Low-Risk, Moderate-Risk, Suspicious, and High-Risk. If there is no traffic on a specific reputation, then the corresponding reputation does not appear in the chart. The circles in this chart are click-able. Clicking on circle filters rest of page data with the selected reputation as filter and this chart is locked until the filter is removed by clicking on **Remove filter on <reputation>**. For example, see the following figure.

Figure 151 *Filter by Reputation*



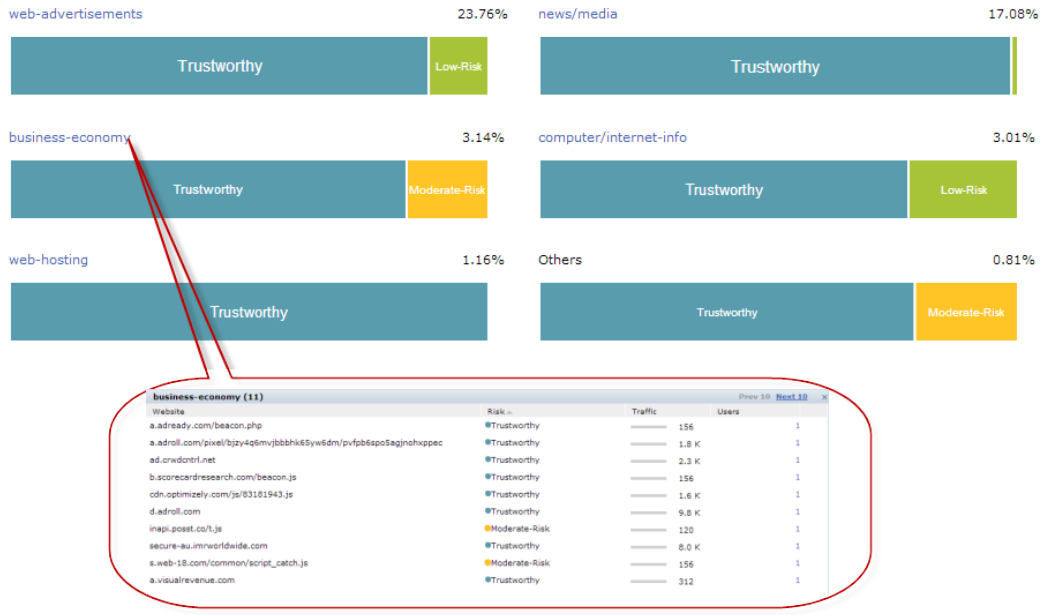
- **Category Views**: A drop-down at the extreme right of reputation traffic lights allows selecting the category view. The view options are Top 9 and Top 6. Top 9 is the default view and displays predefined set of categories that need to be listed in categories by reputation chart. This also list the top 6 or top 9 categories based on traffic usage. The list updates automatically when filters are applied. The following figure shows an for Top 9 category view with reputation chart.

Figure 152 *Category View- Top 9*



- **Details Table:** Click on the web category link above the Category view chart to display the details table as shown in the following figure.

Figure 153 *Category Views and Details*

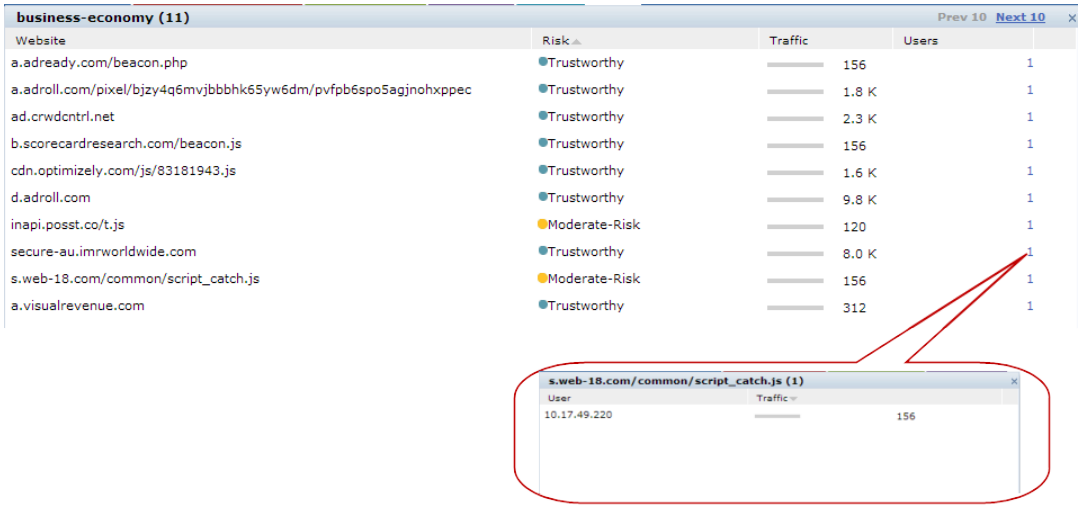


The details table of the selected web category includes the following four columns:

- **Website**: Lists the website
- **Risk**: Reputation score of the website with image presentation
- **Traffic**: Traffic of the website in total traffic of the selected category
- **User**: The number of users using that website

- **User Table:** Click on the number in the **User** column in the details table as shown in the following figure:

Figure 154 *User Table*



The user table includes the following columns:

- **User**: Lists the users of the website
- **Traffic**: Traffic of the user on the website

Web Content Filters

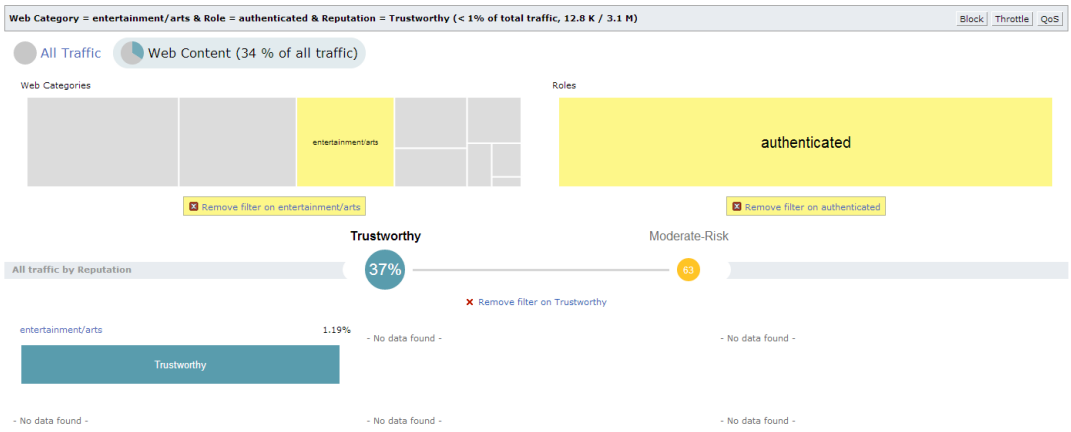
Web content tree chart filter behaves in the same way described in Filters on page 759. Filters can be applied to Web Categories, Roles, and Reputation containers.

Following are the properties of container filters:

- Clicking on any box in the tree chart or reputation traffic light chart will update whole page with the selected box as filter.
- On clicking, the tree chart will freeze that chart and update rest of the page.
- Filter will be applied only to non-freeze chart.
- Reputation chart color won't change upon selection.

The following figure shows an example with multiple filters:

Figure 155 *Multiple Filters*



WebCC Configuration in the WebUI

Configurations of policies from web content dashboard can be done with the help of the following Block/Unblock, Throttle, and QoS Action Buttons. These buttons behave the same way as described in Block/Unblock, Throttle, and QoS Action Buttons.

Block / Unblock:

To permit or deny a rule for global policy or per-role policies for a web category, role, or reputation. To apply a policy, click on a on a web category, role, reputation, or a combination of these three container and click block. Click **OK**. For example, the following two figures show applying a policy on web category filter and on Role + Category + Reputation filter:

Figure 156 *Policy on Web Category Filter*

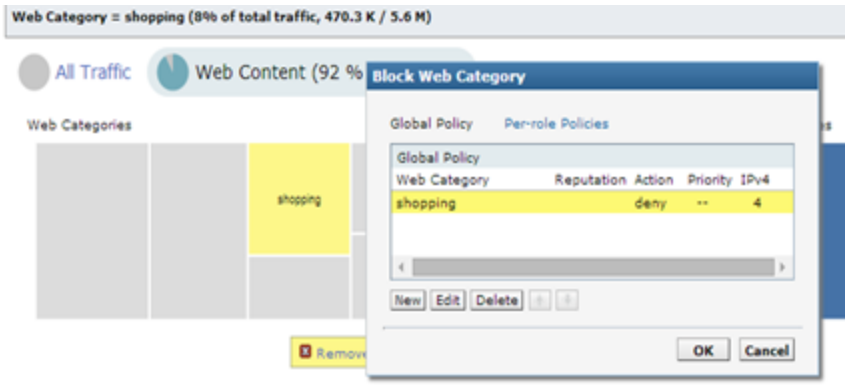
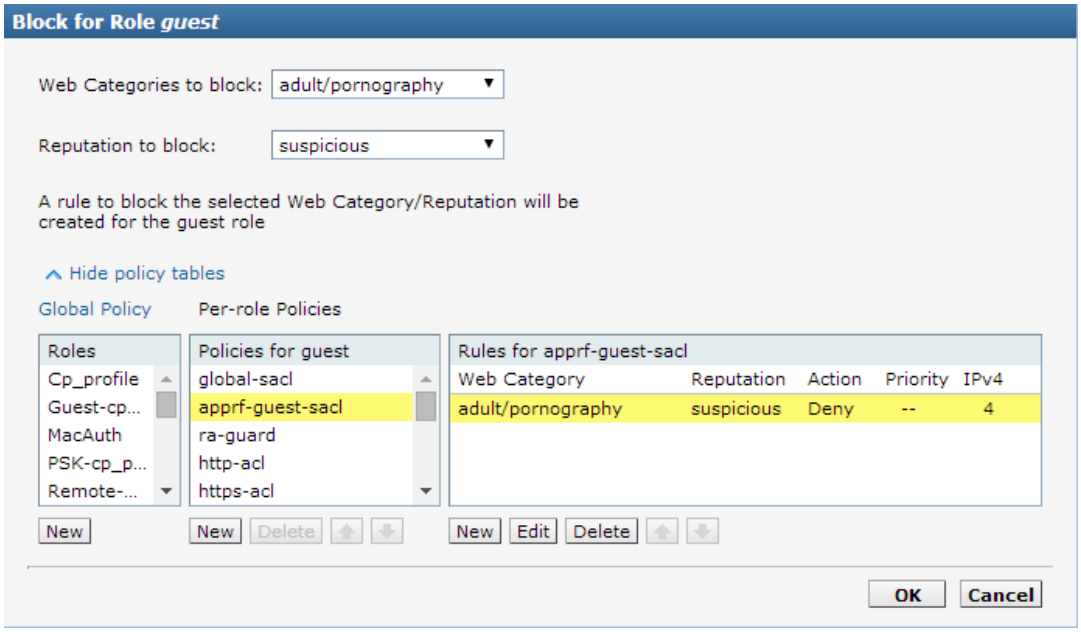


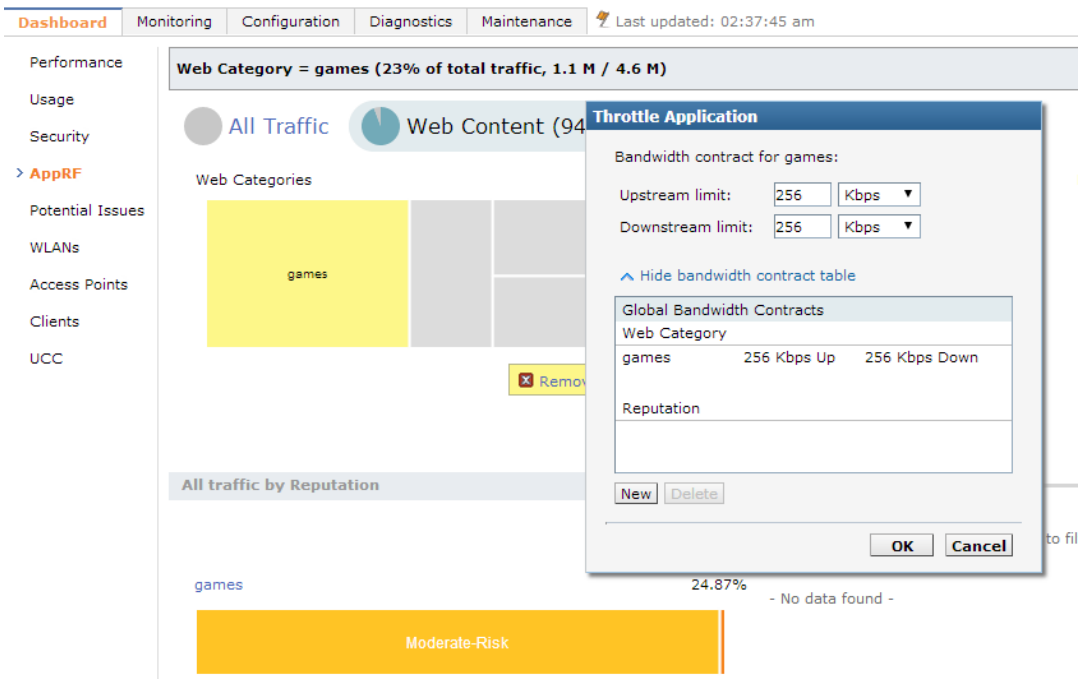
Figure 157 *Policy on Web Category + Reputation + Role Filter*



Throttle

To apply bandwidth contract for a web category, role, or reputation. For example, the following figure shows the throttle applied to a category filter:

Figure 158 *Throttle on Category Filter*



When multiple bandwidth contracts exist, the precedence is as follows:

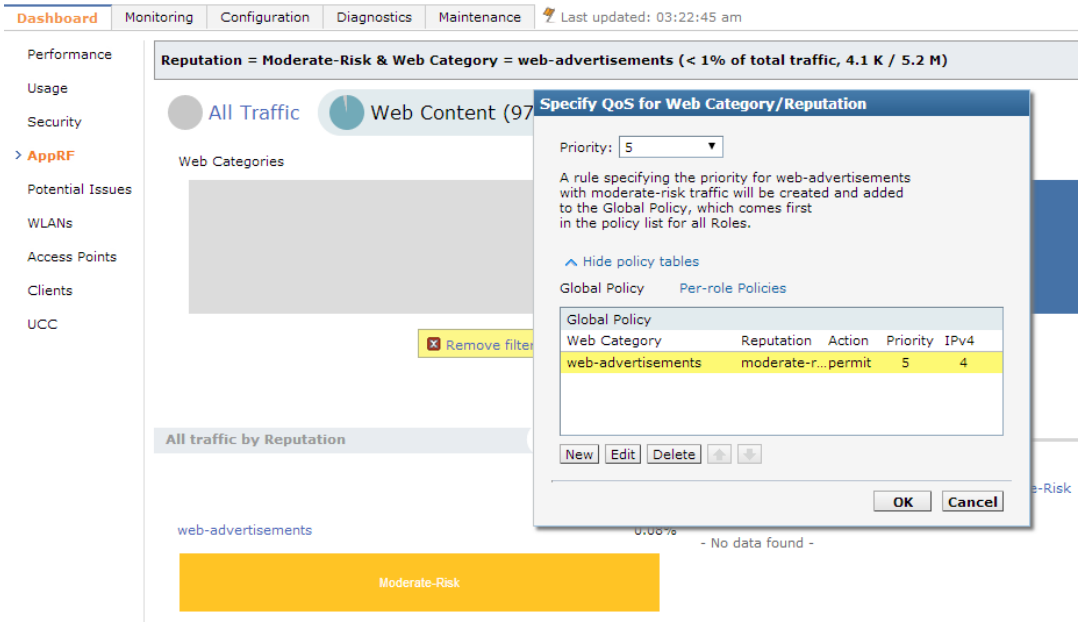
- WebCC Global bandwidth contract
- Application bandwidth exception List
- Application Category bandwidth exception List
- App bandwidth contract
- Application Category bandwidth contract

- Web category bandwidth contract
- Web reputation bandwidth contract
- User bandwidth contract

QoS

To set the priority of the web category and reputation. For example, the following figure shows QoS on category and reputation filter:

Figure 159 *QoS on Web Category + Reputation Filter*



Additionally rules can be added in any of the following combination:

- Rules for Web category only
- Rules for Reputation only
- Rules for Web Category and Web Reputation combination

WebCC Configuration in the CLI

Enabling WebCC

Use the following command to enable WebCC using the CLI:

```
(host) (config) #firewall web-cc
```

Use the following command to configure WebCC per-role using the CLI:

```
(host) (config-role) #web-cc
```

New policy configuration

The new CLI extends the existing policy configuration to take web category or reputation or both. Use the following command to configure a new policy to create ACL rule with web category and reputation:

```
(host) (config-sess-acl) #source destination proto-port/service/app/app-group <name> webcc-category <ctgry> webcc-reputation <score> action [log | mirror | time-range]
```

The following actions are supported when web category/reputation is selected:

- Deny
- Permit

- Blacklist
- Classify-media
- Disable-scanning
- Dot1q-priority
- Log
- Mirror
- Queue
- Time-range
- TOS

Example for WebCC policy configuration is as follows:

```
ip access-list session url-filter
   any any web-cc-category educational-institutions permit
   any any web-cc-reputation suspicious deny
   any any any deny
```

Assuming that webcc categorization was done only for `http` traffic running on TCP 80, the above ACL is converted as follows in datapath for pre-classification ACL scan:

```
ip access-list session url-filter
   any any tcp {80} permit
   any any tcp {80} deny
   any any any deny
```

Post-classification, ACL look-up will have the ACL as follows:

```
ip access-list session url-filter
   any any tcp {80} WebCCCtgID 40 WebCCRep 1-100 permit
   any any tcp {80} WebCCRep 1-100 deny
   any any any deny
```

In case there exists an ACL rule to deny/permit a specific web category but is required to make an exception to allow/deny a specific URL or website, then this can be accomplished by configuring in the following manner:

1. First define a netdestination with one or more URLs to whitelist or blacklist

   ```
   (config) #netdestination search
   (config-dest) #name www.google.com
   (config-dest) #name www.bing.com
   (config-dest) #exit
   ```

2. Apply this netdestination to an ACL

   ```
   (config) #ip access-list session whitelist
   (config-sess-whitelist)#any alias search tcp 80 permit
   (config-sess-whitelist)#any alias search tcp 443 permit
   ```

3. Apply this ACL to an user-role. The position of this ACL should be at the top. However, with global or role-specific default ACLs this wouldn't be possible.

   ```
   (config) #user-role guest2
   (config-role) #access-list session whitelist
   ```



If there a web-cc/app rule that is applicable globally across user-roles, then there is no way to override such behavior. This is a limitation.

WebCC Bandwidth Contract Configuration

With this feature, ArubaOS supports configuring WebCC category and reputation based bandwidth contract configuration/enforcement. This can be enforced globally for all user-roles, or can be enforced per user-role.

Use the following command to apply global WebCC based bandwidth contracts using the CLI:

```
(host) (config) #web-cc global-bandwidth-contract webcc-category/webcc-reputation  <name>
upstream/downstream mbits/kbits <value>
```

Use the following command to apply AAA bandwidth contracts using the CLI:

```
(host) (config) #aaa bandwidth-contract webcc mbits <value>
```

Use the following command to apply role-specific web-cc based bandwidth contracts using the CLI:

```
(host) (config) #user-role webcc
(host) (config-role) #bw-contract webcc-category/webcc-reputation <name> <contract>
upstream/downstream
```

Debugging— The following **show** commands are introduced as part of this feature:

- **show web-cc category all**: Displays all WebCC categories
- **show web-cc reputation:** Displays WebCC reputation
- **show web-cc stats:** Displays the statistics of WebCC module in CP
- **show web-cc status:** Display the status of Web-CC module in CP
- **show web-cc global-bandwidth-contract:** Displays configured WebCC bandwidth contract
- **show datapath web-cc:** Displays md5, web category, reputation, and age for each URL
- **show datapath web-cc counters:** Displays the number of URLs in cache, Classified and Unclassified sessions.
- **show datapath session web-cc:** Displays Internal Flags, Pre Classification ACE Index, and Post Classification ACE Index
- **show gsm debug channel web_cc_info**: Lists md5, Category, and Reputation for each URL. GSM entries are populated as and when URL cache entry is learned, and it is used for reporting the actual URLs being associated with user session entries.

The following **clear** command are introduced as part of this feature:

- **clear web-cc cache <md5_1> <md5_2>** : Clears the WebCC cache entry from both data plane and GSM.
- **clear web-cc stats**: Clears all WebCC statistics.
- **clear datapath web-cc counters**: Clears configuration values and statistics in the WebCC datapath module.

AirGroup

The **Dashboard** tab of the controller WebUI contains an **AirGroup** link that displays the information about AirGroup clients and servers. By default, these tables contain information for all active AirGroup clients and servers. You can filter the information in these tables by clicking the filter icon on any column heading and entering a string in the filter field.

The **Dashboard>AirGroup** window contains the following information for AirGroup Users and Servers:

Table 156: *AirGroup Monitoring Information*

Column	
AirGroup Users	
Host Name	Host name of the AirGroup server
User Name	User name given to a client that completed 802.1X authentication

Column	
IP address	Device IP address
Role	Role assigned to the device's user
AP Name	Name of the AP to which the device is associated
VLAN ID	ID of the VLAN to which the device is assigned
Group(s)	Displays the Group of the AirGroup user
AirGroup Type	Displays the type of the device
Wired/Wireless	Type of connection between the device and the LAN
AirGroup Servers	
Host Name	Host name of the AirGroup server
Service	Service(s) running on the server
IP address	AirGroup Server's IP address
MAC	AirGroup Server's MAC address
Role	Role assigned to the AirGroup server
Wired/Wireless	Type of connection between the device and the LAN
AP Name	Name of the AP to which the device is associated
VLAN ID	ID of the VLAN to which the server is assigned
Group(s)	Displays the Group of the AirGroup user
AirGroup Type	Displays the type of the device

For more information on the AirGroup feature, see AirGroup on page 940

UCC

The Unified Communication and Collaboration (UCC) Dashboard Aggregated Display shows an aggregated view of the UCC calls made in the controller. The administrator can see a top level view of the call quality assessment, and further drill down into a specific view based on the analysis required.

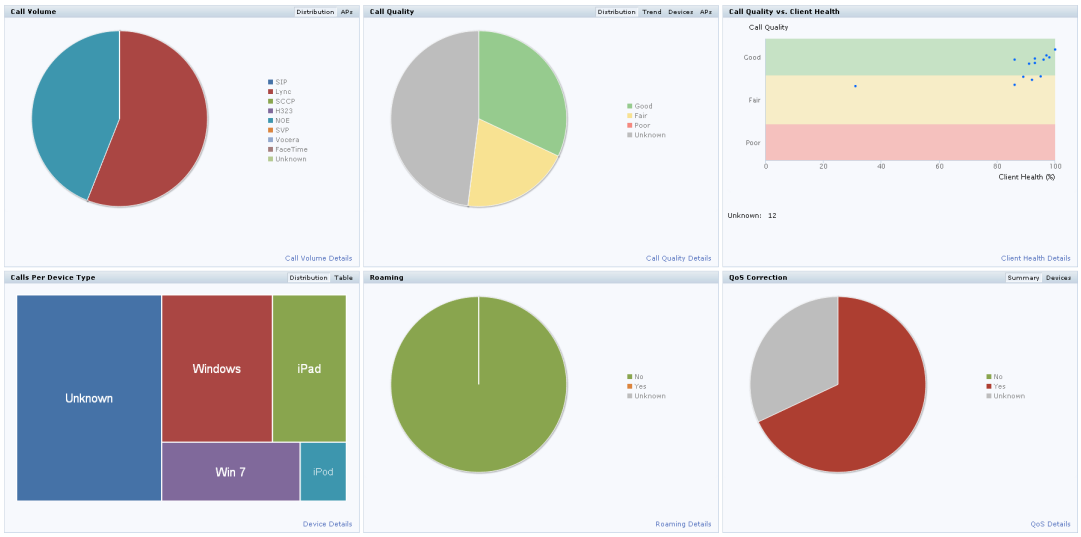


The UCC feature requires the PEFNG license.

Chart View

A new **UCC** tab is introduced under the **Dashboard** tab. Navigate to the **Dashboard > UCC** page to view UCC dashboard. Clicking the **UCC** hyperlink displays the following characteristics (in graphical format) of the UCC deployment.

Figure 160 *UCC Dashboard*



- **Call Volume** – This graph displays the total number of calls made based on the UCC application type. For example, SIP, Lync, SCCP, H.323, NOE, SVP, VOCERA, and FaceTime.
- **Call Quality** – This graph displays the number of UCC calls categorized by the following call quality:
  - Good
  - Fair
  - Poor
  - Unknown
- **Call Quality vs. Client Health** - This graph displays the co-relation between the VoIP call quality and the VoIP client health of every UCC call.
- **Calls Per Device Type** – This graph displays the calls made per device type. For example, Windows 7, Mac OS X, iPhone, or Android.
- **Roaming** – Roaming status of UCC clients. The status can be:
  - No – Number of calls where the client did not roam to a new AP.
  - Yes - Number of calls where the client has roamed to a new AP.
  - Unknown - Unknown calls.
- **QoS Correction** – If the DSCP value of the Real-time Transport Protocol (RTP) packets sent by the client differs from the recommended QoS setting, the call is classified as QoS Corrected. This graph displays the number of UCC calls where the controller has corrected the DSCP QoS value for such calls. The QoS correction is categorized as:
  - No – No UCC QoS call correction.
  - Yes – DSCP QoS value corrected by the controller.
  - Unknown – Unknown calls.

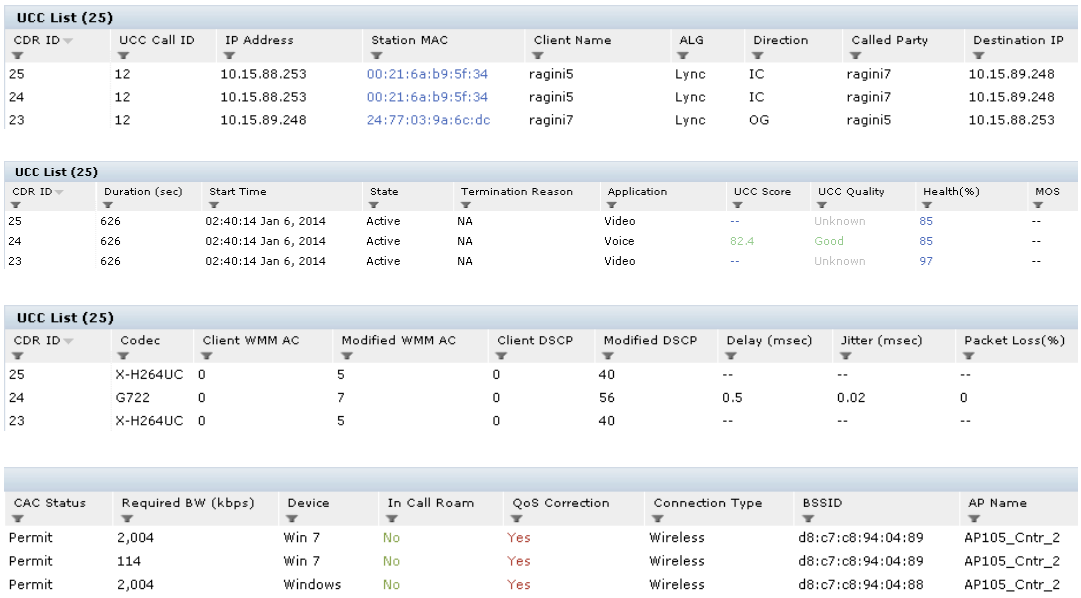
Details View

Navigate to the **Dashboard > UCC** page. To display an aggregated list of all the UCC call data metrics in the controller, click any of the following hyperlinks:

- Call Volume Details
- Call Quality Details
- Client Health Details
- Device Details
- Roaming Details
- QoS Details

Figure 161 displays an aggregated list of all the UCC call data metrics in the controller.

Figure 161 *UCC List*



UCC List (25)

CDR ID	UCC Call ID	IP Address	Station MAC	Client Name	ALG	Direction	Called Party	Destination IP
25	12	10.15.88.253	00:21:6a:b9:5f:34	ragini5	Lync	IC	ragini7	10.15.89.248
24	12	10.15.88.253	00:21:6a:b9:5f:34	ragini5	Lync	IC	ragini7	10.15.89.248
23	12	10.15.89.248	24:77:03:9a:6c:dc	ragini7	Lync	OG	ragini5	10.15.88.253

UCC List (25)

CDR ID	Duration (sec)	Start Time	State	Termination Reason	Application	UCC Score	UCC Quality	Health(%)	MOS
25	626	02:40:14 Jan 6, 2014	Active	NA	Video	--	Unknown	85	--
24	626	02:40:14 Jan 6, 2014	Active	NA	Voice	82.4	Good	85	--
23	626	02:40:14 Jan 6, 2014	Active	NA	Video	--	Unknown	97	--

UCC List (25)

CDR ID	Codec	Client WMM AC	Modified WMM AC	Client DSCP	Modified DSCP	Delay (msec)	Jitter (msec)	Packet Loss(%)
25	X-H264UC	0	5	0	40	--	--	--
24	G722	0	7	0	56	0.5	0.02	0
23	X-H264UC	0	5	0	40	--	--	--

CAC Status	Required BW (kbps)	Device	In Call Roam	QoS Correction	Connection Type	BSSID	AP Name
Permit	2,004	Win 7	No	Yes	Wireless	d8:c7:c8:94:04:89	AP105_Cntr_2
Permit	114	Win 7	No	Yes	Wireless	d8:c7:c8:94:04:89	AP105_Cntr_2
Permit	2,004	Windows	No	Yes	Wireless	d8:c7:c8:94:04:88	AP105_Cntr_2

This chapter describes management access and tasks for a user-centric network and includes the following topics:

- Configuring Certificate Authentication for WebUI Access on page 778
- Secure Shell (SSH) on page 779
- Enabling RADIUS Server Authentication on page 780
- Connecting to an AirWave Server on page 785
- Custom Certificate Support for RAP on page 786
- Implementing a Specific Management Password Policy on page 788
- Configuring AP Image Preload on page 790
- Configuring Centralized Image Upgrades
- Managing Certificates on page 795
- Configuring SNMP on page 801
- Enabling Capacity Alerts on page 802
- Configuring Logging on page 804
- Enabling Guest Provisioning on page 806
- Managing Files on the Controller on page 821
- Setting the System Clock on page 824
- ClearPass Profiling with IF-MAP on page 826
- Whitelist Synchronization on page 827
- Downloadable Regulatory Table on page 828

Configuring Certificate Authentication for WebUI Access

The controller supports client certificate authentication for users accessing the controller using the WebUI. (The default is for username/password authentication.) You can use client certificate authentication only, or client certificate authentication with username/password (if certificate authentication fails, the user can log in with a configured username and password).



Each controller can support a maximum of ten management users.

To use client certificate authentication, you must do the following:

1. Obtain a client certificate and import the certificate into the controller. Obtaining and importing a client certificate is described in Managing Certificates on page 795.
2. Configure certificate authentication for WebUI management. You can optionally also select username/password authentication.
3. Configure a user with a management role. Specify the client certificate for authentication of the user.

In the WebUI

1. Navigate to the **Configuration > Management > General** page.

2. Under WebUI Management Authentication Method, select Client Certificate. You can select Username and Password as well; in this case, the user is prompted to manually enter the username and password only if the client certificate is invalid.
3. Select the server certificate to be used for this service.
4. Click **Apply**.
5. To configure the management user, navigate to the **Configuration > Management > Administration** page.
   a. Under Management Users, click **Add**.
   b. Select Certificate Management.
   c. Select WebUI Certificate.
   d. Enter the username.
   e. Select the user role assigned to the user upon validation of the client certificate
   f. Enter the serial number for the client certificate.
   g. Select the name of the CA that issued the client certificate.
   h. Click **Apply**.

In the CLI

```
web-server
   mgmt-auth certificate
   switch-cert <certificate>
mgmt-user webui-cacert <ca> serial <number> <username> < role>
```

Secure Shell (SSH)

SSH is enabled by default in ArubaOS, and thus lets you log in using a username and password. You can enable SSH login by using public key authentication while leaving username/password authentication enabled, or you may disable the username/password authentication and leave only the public key authentication enabled. In the FIPS mode of operation, SSH is pre-configured to only use Diffie-Hellman Group 14 with AES-CBC-128 and AES-CBC-256 and HMAC-SHA1/HMAC-SHA1-96. These settings are not configurable.

When you import an X.509 client certificate into the controller, the certificate is converted to SSH-RSA keys. When you enable public key authentication for SSH, the controller validates the client's credentials with the imported public keys. You can specify public key authentication only, or public key authentication with username/password (if the public key authentication fails, the user can login with a configured username and password).

Enabling Public Key Authentication

The controller allows public key authentication of users accessing the controller using SSH. (The default is for username/password authentication.)

To use public key authentication, you must do the following:

1. Import the X.509 client certificate into the controller using the WebUI, as described in Importing Certificates on page 798
2. Configure SSH for client public key authentication. You can optionally also select username/password authentication.
3. Configure the username, role and client certificate.

In the WebUI

1. Navigate to the **Configuration > Management > General** page.
2. Under SSH (Secure Shell) Authentication Method, select Client Public Key. You can optionally select Username/Password to use both username/password and public key authentication for SSH access.
3. Click **Apply**.
4. To configure the user, navigate to the **Configuration > Management > Administration** page.
   a. Under Management Users, click **Add**.
   b. Select Certificate Management.
   c. Select SSH Public Key.



ArubaOS recommends that the username and role for SSH be the same as for the WebUI Certificate. You can optionally use the checkbox to copy the username and role from the Web Certificate section to the SSH Public Key section.

   d. Enter the username.
   e. Select the management role assigned to the user upon validation of the client certificate.
   f. Select the client certificate.
   g. Click **Apply**.

In the CLI

```
ssh mgmt-auth public-key [username/password]
mgmt-user ssh-pubkey client-cert <certificate> <username> <role>
```

Enabling RADIUS Server Authentication

This section include many different types of RADIUS server configuration and related procedures.

Configuring RADIUS Server Username and Password Authentication

In this example, an external RADIUS server is used to authenticate management users. Upon authentication, users are assigned the default role root.

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **RADIUS Server** to display the Radius Server List.
   a. To configure a RADIUS server, enter the name for the server (for example, rad1) and click **Add**.
   b. Select the name to configure server parameters, such as IP address. Select the **Mode** checkbox to activate the server.
   c. Click **Apply**.
3. Select **Server Group** to display the Server Group list.
   a. Enter the name of the new server group (for example, corp_rad) and click **Add**.
   b. Select the name to configure the server group.
   c. Under Servers, click **New** to add a server to the group.
   d. Select a server from the drop-down menu and click **Add Server**.
   e. Click **Apply**.
4. Navigate to the **Configuration > Management > Administration** page.

a. Under Management Authentication Servers, select a management role (for example, root) for the Default Role.
b. Select (check) Mode.
c. For Server Group, select the server group that you just configured.
d. Click **Apply**.

In the CLI

```
aaa authentication-server radius rad1
  host <ipaddr>
  enable

aaa server-group corp_rad
  auth-server rad1

aaa authentication mgmt
  default-role root
  enable
  server-group corp_rad
```

Configuring RADIUS Server Authentication with VSA

In this scenario, an external RADIUS server authenticates management users and returns to the controller the Dell vendor-specific attribute (VSA) called Dell-Admin-Role that contains the name of the management role for the user. The authenticated user is placed into the management role specified by the VSA.

The controller configuration is identical to the Configuring RADIUS Server Username and Password Authentication on page 780. The only difference is the configuration of the VSA on the RADIUS server. Ensure that the value of the VSA returned by the RADIUS server is one of the predefined management roles. Otherwise, the user will have *no* access to the controller.

Configuring RADIUS Server Authentication with Server Derivation Rule



Dell controllers do not make use of any returned attributes from a TACACS+ server.

A RADIUS server can return to the controller a standard RADIUS attribute that contains one of the following values:

- The name of the management role for the user
- A value from which a management role can be derived

For either situation, configure a server-derivation rule for the server group.

In the following example, the RADIUS server returns the attribute Class to the controller. The value of the attribute can be either “root” or “network-operations” depending upon the user; the returned value is the role granted to the user.



Ensure that the value of the attribute returned by the RADIUS server is one of the predefined management roles. Otherwise, the management user will not be granted access to the controller.

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **RADIUS Server** to display the Radius Server List.
   a. To configure a RADIUS server, enter the name for the server (for example, rad1) and click **Add**.

b. Select the name to configure server parameters, such as IP address. Select the **Mode** checkbox to activate the server.
c. Click **Apply**.

3. Select **Server Group** to display the Server Group list.
   a. Enter the name of the new server group (for example, corp_rad) and click **Add**.
   b. Select the name to configure the server group.
   c. Under Servers, click **New** to add a server to the group.
   d. Select a server from the drop-down menu and click **Add Server**.
   e. Under Server Rules, click **New** to add a server rule.
   f. For Condition, select **Class** from the scrolling list. Select **value-of** from the drop-down menu. Select **Set Role** from the drop-down menu.
   g. Click **Add**.
   h. Click **Apply**.
4. Navigate to the **Configuration > Management > Administration** page.
   a. Under Management Authentication Servers, select a management role (for example, read-only) for the Default Role.
   b. Select (check) Mode.
   c. For Server Group, select the server group that you just configured.
   d. Click **Apply**.

In the CLI

```
aaa authentication-server radius rad1
  host <ipaddr>
  enable

aaa server-group corp_rad
  auth-server rad1
  set role condition Class value-of

aaa authentication mgmt
  default-role read-only
  enable
  server-group corp_rad
```

In the following example, the RADIUS server returns the attribute Class to the controller; the value of this attribute can be “it”, in which case, the user is granted the root role. If the value of the Class attribute is anything else, the user is granted the default read-only role.

Configuring a set-value server-derivation rule

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **RADIUS Server** to display the Radius Server List.
   a. To configure a RADIUS server, enter the name for the server (for example, rad1) and click **Add**.
   b. Select the name to configure server parameters, such as IP address. Select the **Mode** checkbox to activate the server.
   c. Click **Apply**.
3. Select **Server Group** to display the Server Group list.
   a. Enter the name of the new server group (for example, corp_rad) and click **Add**.

b. Select the name to configure the server group.
c. Under Servers, click **New** to add a server to the group.
d. Select a server from the drop-down menu and click **Add Server**.
e. Under Server Rules, click **New** to add a server rule.
f. For Condition, select **Class** from the scrolling list. Select **equals** from the drop-down menu. Enter **it**. Select **Set Role** from the drop-down menu. For Value, select **root** from the drop-down menu.
g. Click **Add**.
h. Click **Apply**.

4. Navigate to the **Configuration > Management > Administration** page.
   a. Under Management Authentication Servers, select a management role (for example, read-only) for the Default Role.
   b. Select (check) Mode.
   c. For Server Group, select the server group that you just configured.
   d. Click **Apply**.

In the CLI

```
aaa authentication-server radius rad1
  host <ipaddr>
  enable

aaa server-group corp_rad
  auth-server rad1
  set role condition Class equals it set-value root

aaa authentication mgmt
  default-role read-only
  enable
  server-group corp_rad
```

For more information about configuring server-derivation rules, see Configuring Server-Derivation Rules on page 242.

Disabling Authentication of Local Management User Accounts

You can disable authentication of management user accounts in local switches if the configured authentication server(s) (RADIUS or TACACS+) are not available.

You can disable authentication of management users based on the results returned by the authentication server. When configured, locally-defined management accounts (for example, admin) are not allowed to log in if the server(s) are reachable and the user entry is not found in the authentication server. In this situation, if the RADIUS or TACACS+ server is unreachable, meaning it does not receive a response during authentication, or fails to authenticate a user because of a timeout, local authentication is used and you can log in with a locally-defined management account.

In the WebUI

1. Navigate to the **Configuration > Management > Administration** page.
2. Under Management Authentication Servers, uncheck the **Local Authentication Mode** checkbox.
3. Click **Apply**.

In the CLI

```
mgmt-user localauth-disable
```

Verifying the configuration

To verify if authentication of local management user accounts is enabled or disabled, use the following command:

```
show mgmt-user local-authentication-mode
```

Resetting the Admin or Enable Password

This section describes how to reset the password for the default administrator user account (**admin**) on the controller. Use this procedure if the administrator user account password is lost or forgotten.

1. Connect a local console to the serial port on the controller.
2. From the console, login in the controller using the username **password** and the password **forgetme!**.
3. Enter enable mode by typing in **enable**, followed by the password **enable**.
4. Enter configuration mode by typing in **configure terminal**.
5. To configure the administrator user account, enter **mgmt-user admin root**. Enter a new password for this account. Retype the same password to confirm.
6. Exit from the configuration mode, enable mode, and user mode.

This procedure also resets the enable mode password to **enable**. If you have defined a management user password policy, make sure that the new password conforms to this policy. For details, see Implementing a Specific Management Password Policy on page 788.

Figure 162 is an example of how to reset the password. The commands in bold type are what you enter.

Figure 162 *Resetting the Password*

```
(host)
User: password
Password: forgetme!
(host) >enable
Password: enable
(host) #configure terminal
Enter Configuration commands, one per line. End with CNTL/Z

(host) (config) #mgmt-user admin root
Password: ******
Re-Type password: ******
(host) (config) #exit
(host) #exit
(host) >exit
```

After you reset the administrator user account and password, you can login to the controller and reconfigure the enable mode password. To do this, enter configuration mode and type the **enable secret** command. You are prompted to enter a new password and retype it to confirm. Save the configuration by entering **write memory**.

Figure 163 details an example reconfigure the enable mode password. Again, the command you enter displays in bold type.

Figure 163 *Reconfigure the enable mode password*

```
User: admin
Password: ******
(host) >enable
Password: ******
(host) #configure terminal
Enter Configuration commands, one per line. End with CNTL/Z

(host) (config) #enable secret
Password: ******
Re-Type password: ******
(host) (config) #write memory
```

Bypassing the Enable Password Prompt

The bypass enable feature lets you bypass the enable password prompt and go directly to the privileged commands (config mode) after logging on to the controller. This is useful if you want to avoid changing the enable password due to company policy.

Use the `enable bypass` CLI command to bypass the enable prompt an go directly to the privileged commands (config mode). Use the `no enable bypass` CLI command to restore the enable password prompt.

Setting an Administrator Session Timeout

You can configure the number of seconds after which an Administrator's WebUI or CLI session times out.

In the WebUI

To define a timeout interval for a WebUI session, use the command:

```
web-server sessiontimeout <session-timeout>
```

In the above command, `<session-timeout>` can be any number of seconds from 30 to 3600, inclusive.

In the CLI

To define a timeout interval for a CLI session, use the command:

```
loginsession timeout <value>
```

In the above command, `<val>` can be any number of minutes from 5 to 60 or seconds from 1 to 3600, inclusive. You can also specify a timeout value of 0 to disable CLI session timeouts.

Connecting to an AirWave Server

AirWave is a powerful and easy-to-use network operations system that manages Dell W-Series wireless, wired and remote access networks, as well as wireless and wired infrastructures from Dell and a wide range of third-party manufacturers.

Controllers running ArubaOS 6.3 and later can use the AirWave wizard in the **Configuration > Wizards > AirWave** section of the controller WebUI to quickly and easily connect the controller to an AirWave server. The following checklist lists the information you will need to use this wizard. Determine each of these values for your deployment and AirWave server before you start the wizard process.

Table 157: *AirWave Wizard Checklist*

Information	Description	My Values
AirWave IP address	IP address of the AirWave server.	
SNMP version	Specify if the controller and AirWave serer should communicate using SNMP v2 or SNMPv3. SNMPv3 communications between a controller and an AirWave server use SHA authentication and AES encryption.	
For SNMPv2	If you select SNMPv2, you must enter an **SNMP community** string.	
For SNMPv3	If you select SNMPv3, you must enter values for the following parameters: • **User name** : A string representing the name of the SNMP user. • **Authentication password**: Authentication key for use with the SHA authentication protocol. • **privacy password**: Privacy key for encrypted messages. • **NTP serve**r: If the controller is not already configured to use an NTP server, enter the IP address of an NTP server.	
Syslog	Syslog messages are disabled by default. Use the Syslog section of the wizard to enable syslog messages, and define the syslog category, syslog facility levels (local0-local7) and syslog severity levels (debug-emergency) for messages from the controller. By default, AirWave syslog messages sent at the **error** severity level. The possible syslog categories are as follows: • ap-debug • arm-user-debug • network • security • system • user • user-debug • wireless	

Custom Certificate Support for RAP

As Suite-B mandates using the AES-GCM encryption and ECDSA certificates for security, this feature allows you to upload custom RSA and ECDSA certificates to a RAP. This allows custom certificates to be used for IKEv2 negotiation which establishes a tunnel between the RAP and the controller. Feature support includes the ability to:

- Upload a single CA certificate and RAP certificate which have either elliptical crypto key parameters with ECDSA or RSA parameters for signing and verification.
- Store the certificate in the flash of the RAP
- Delete certificates

- Generate a CSR paired with a private key generation for the RAP. The private key is stored in the flash and the CSR can be exported out of the RAP to get it signed by the CA.

If there is a custom certificate present in the flash when rebooting, this feature creates a suite B tunnel with the controller if the certificates uploaded are using EC algorithms. Otherwise it creates a tunnel using standard RAP IPSec parameters.

Suite-B Support for ECDSA Certificate

If a custom ECDSA certificate is present in the flash of a certificate-based RAP, it is automatically designated as a Suite-B RAP. On the controller side, tunnel creation uses the server certificate as a default VPN server certificate.

Administering Suite-B support for a RAP includes these steps which are described in the following sections:

1. Setting the Default Server Certificate
2. Import a custom certificate
3. Generate a Certificate Signing Request (CSR)
4. Upload the certificate

Setting the Default Server Certificate

In the CLI

To set the default server certificate that is presented to the RAP as the default VPN server certificate:

```
(host) (config) #crypto-local isakmp server-certificate
<server_certificate_name>
```

To add the CA certificate to verify the RAP certificate:

```
(host) (config) #crypto-local isakmp ca-certificate <trusted CA>
```

Importing a Custom Certificate

Certificates can only be imported to the controller using the WebUI.

In the WebUI

1. Navigate to **Configuration > Management > Certificates** and upload the certificate.
2. To use imported certificates to create a tunnel, navigate to **Configuration > Advanced Services > Emulate VPN Services**.

Generating a CSR

The RAP console page allows you to generate a CSR. This is done through a private key which can be generated and saved to the RAP flash. A corresponding CSR is exported so it can be signed by the required CA to use as the RAP certificate. This RAP certificate can then be uploaded using the Upload button on the RAP Console page.

The subject of the RAP certificate needs to be the MAC address of the RAP, and nothing more. Note that this is case insensitive.

If you create a CSR on the RAP and then have a certificate issued by a CA, you must have the certificate in PEM format before uploading it to the RAP.

Uploading the Certificate



When using the “rapconsole.dell.com” page on a bridge/split-tunnel RAP to manage certificates on the RAP, a blank page or a page that does not have the Certificates tabs on it may display. The RAP provisioning page that is standard on the RAP may conflict with the “rapconsole” page and thus confuse the browser. If this occurs, clear your browser cache first or use two different browsers.

The Upload button on the RAP console page that lets you upload the certificates to the RAP flash. The certificate needs to be in PEM format and uploading the RAP certificate requires that the corresponding private key is present in the RAP flash. Or, use the PKCS12 bundle where the chain includes the RAP private key with the RAP and CA certificates are optionally password protected.

Implementing a Specific Management Password Policy

By default, the password for a new management user has no requirements other than a minimum length of 6 alphanumeric or special characters. However, if your company enforces a best practices password policy for management users with root access to network equipment, you may want to configure a password policy that sets requirements for management user passwords.

Defining a Management Password Policy

To define specific management password policy settings through the WebUI or the CLI, complete the following steps:

In the WebUI

1. Navigate to **Configuration>All Profiles**.
2. Expand **Other Profiles**.
3. Select **Mgmt Password Policy**.
4. Configure the settings described in Table 158.

Table 158: *Management Password Policy Settings*

Parameter	Description
Enable Password Policy	Select this checkbox to enable the password management policy. The password policy will not be enforced until this checkbox is selected.
Minimum password length required	The minimum number of characters required for a management user password Range: 6-64 characters. Default: 6.
Minimum number of Upper Case characters	The minimum number of uppercase characters required in a management user password. Range: 0-10 characters. By default, there is no requirement for uppercase letters in a password, and the parameter has a default value of 0.
Minimum number of Lower Case characters	The minimum number of lowercase characters required in a management user password. Range: 0-10 characters. By default, there is no requirement for lowercase letters in a password, and the parameter has a default value of 0.
Minimum number of Digits	The minimum number of numeric digits required in a management user password. Range: 0-10 digits. By default, there is no requirement for numerical digits in a password, and the parameter has a default value of 0.

Table 158: *Management Password Policy Settings*

Parameter	Description
Minimum number of Special characters (!, @, #, $, %, ^, &, *, <, >, {, }, [,], :, ., comma, |, +, ~, `)	The minimum number of special characters. Range: 0-10 characters.
Username or Reverse of username NOT in Password	When you select this checkbox, the password cannot be the management users' current username or the username spelled backwards.
Maximum consecutive character repeats	The maximum number of consecutive repeating characters allowed in a management user password. Range: 0-10 characters. By default, there is no limitation on the numbers of character that can repeat within a password, and the parameter has a default value of 0 characters.
Maximum Number of failed attempts in 3 minute window to lockout user	The number of failed attempts within a 3 minute window that causes the user to be locked out for the period of time specified by the **Time duration to lockout the user upon crossing the "lock-out" threshold** parameter. Range: 0-10 attempts. By default, the password lockout feature is disabled, and the default value of this parameter is 0 attempts.
Time duration to lock out the user upon crossing the "lock-out" threshold	The duration in time that locks out the user upon crossing the lock out threshold. Range: 0-60 in minutes.

5. Click **Apply** to save your settings.

In the CLI

```
aaa password-policy mgmt
   enable
   no
   password-lock-out
   password-lock-out-time
   password-max-character-repeat.
   password-min-digit
   password-min-length
   password-min-lowercase-characters
   password-min-special-character
   password-min-uppercase-characters
   password-not-username
```

Management Authentication Profile Parameters

Table 159 describes configuration parameters on the Management Authentication profile page.



In the CLI, you configure these options with the **aaa authentication mgmt** and **aaa-server-group** commands.

Table 159: *Management Authentication Profile Parameters*

Parameter	Description
Enable	Enables authentication for administrative users.
Default Role	Select a predefined management role to assign to authenticated administrative users:
Root	Default superuser role
guest-provisioning	Guest provisioning role
location-api-mgmt	Location API role
network-operations	Network operations role
no-access	No commands are accessible for this role
read-only	Read-only role
no access	Negates any configured parameter.
Server Group	Name of the group of servers used to authenticate administrative users. See the CLI command **`aaa-server-group`**, in the *CLI Command Reference Guide* for more information.

Configuring AP Image Preload

The AP image preload feature minimizes the downtime required for a controller upgrade by allowing the APs associated to that controller to download the new images before the controller actually starts running the new version.



This feature is supported only on the W-3600, W-6000M3, and W-7200 Series controllers.

This feature allows you to select the maximum number of APs that are allowed to preload the new software image at any one time, thereby reducing the possibility that the controller may get overloaded or that network traffic may be impacted by all APs on the controller attempting to download a new image at once.

APs can continue normal operation while they are downloading their new software version. When the download completes, the AP sends a message to the controller, informing it that the AP has either successfully downloaded the new software version, or that the preload has failed for some reason. If the download fails, the AP will retry the download after a brief waiting period.

You can allow every AP on a controller to preload a new software version, or also create a custom list of AP groups or individual APs that can use this feature. If a new AP associates to the controller while the AP image download feature is active, the controller will check that AP's name and group to see if it appears in the preload list. If an AP is on the list, (and does not already have the specified image in its Flash memory) that AP will start preloading its image.

Enable and Configure AP Image Preload

Use the following procedures to enable and configure the AP Image Preload feature on the W-3600, W-6000M3, and W-7200 Series controllers using the WebUI or CLI.

In the WebUI

1. Navigate to **Maintenance > WLAN > Preload AP Image**. If this feature has not yet been enabled, the window will display the message "AP Image Preload status is Inactive. Click *here to activate AP Image Preload*." Click the link in the warning message to enable this feature and display the AP Image Preload settings.
2. Configure the settings described in the table below, then click **Apply** to save your changes.

Table 160: *AP Image Preload Settings*

Setting	Description
AP Image Preload	Select **Enable** to enable this feature, or **Disable** to disable AP image preload. AP image preload is disabled by default. **NOTE:** This feature can also be enabled and disabled with its current configuration settings in the **Maintenance > Controller > Image Management** window.
Partition	Select the controller partition from which the APs should download their images. By default, the APs will preload images from the controller's default boot partition.
Software Version	This field shows the image on the partition that will be preloaded onto eligible APs, and is not editable.
Maximum Number of Simultaneous downloads	Specify the maximum number of APs that can simultaneously download their image from the controller. A higher number will decrease the time it takes for many APs to preload their new image, but will increase the workload on the controller.
APs to Preload	In this field, select All APsif you want to preload images on all registered APs that are eligible for preload and that support this feature, or select Specific APs to preload images on a list of selected APs. If you selected **Specific APs**, you must create a list of APs allowed to preload images. You can preload images to a group of APs, or specify APs that can use this feature by identifying those APs by AP name. To preload images to a group of APs: 1. In the **AP Groups** field, click **Add**. 2. Type the AP Group Name, or select the AP Group from the list. 3. Click **OK**. (To remove AP groups from this list of APs using this feature, select an AP group name in the list, then click **Delete**.) To preload images to APs with specific name: 1. In the **AP Names** field, click Add. 2. Type the AP Name, or select the AP Name from the list 3. Click **OK**. (To remove an AP from this list of APs using this feature, select an AP name in the list, then click **Delete**.)

In the CLI

To configure the AP image preload feature using the command-line interface, enter the following commands in **enable** mode.

```
ap image-preload
   activate all-aps|specific-aps
   add {ap-group <ap-group> | ap-name <ap-name>}
   cancel
   clear-all
   delete {ap-group <ap-group> | ap-name <ap-name>}
   [partition <part-num>]
   [max-downloads <max-downloads>]
```

The command **ap image preload clear-all** deletes all AP groups and AP names from the list of APs eligible for preloading. This command may be executed either before or after preloading is activated. If it is executed *after* preloading has already been activated, any APs waiting to preload the new software version will be removed

from the list. APs that have already begun the preloading process will continue to download their image and will not be affected.

The **ap image-preload cancel** command deletes all AP groups and AP names from the list of APs eligible for preloading and cancels the preloading process for any APs on the list that have already begun to download the new image. This command then disables the image preload feature.

View AP Preload Status

You can monitor the current preload status of APs using the image preload feature using the **show ap image-preload-status** and **show ap image-preload-status-summary** commands in the command-line interface, or in the **Maintenance > WLAN > Preload AP Image** window in the WebUI.

The output of the **show ap image-preload-status** CLI command and the **AP Image Preload Status** and **AP Image Preload Status Summary** tables in the WebUI contain the following information:

Table 161: *AP Image Preload Status Settings*

Column	Description
AP Image Preload State/Count	These two columns list the different possible preload states for APs eligible to preload a new software image, and the total number of APs in each state. • **Preloaded**: Number of APs that have finished preloaded a new software image. • **Preloading**: Number of APs that are currently downloading the new image. • **Waiting**: Number of APs that are waiting to start preloading the new image from the controller.
Count	This column lists the number of eligible APs currently in each preload state.
AP Name	Name of an AP eligible to preload a new software image.
AP Group	AP group of an AP eligible to preload a new software image.
AP IP	IP address of the AP.
AP Type	AP model type.
Preload State	Current preload state for the AP • **Preloaded**: The AP is finished preloading a new software image. • **Preloading**: The AP is currently downloading the new image. • **Waiting**: The AP is waiting to start preloading the new image from the controller.
Start Time	Time the AP starting preloading an image.
End Time	Time the AP completed the image preload.
Failure Count	Number of times that the AP failed to preload the new image.
Failure Reason	In the event of an image preload failure, this column will display the reason that the image download failed.

Configuring Centralized Image Upgrades

The centralized image upgrade feature introduced in ArubaOS 6.3 allows the master controller to automatically upgrade its associated local controllers by sending an image from a image server to one or more local controllers. If your master controller supports different local controller models, you can upload different image

types to the server, and the centralized image upgrade feature will send the local controller only the type of image that controller supports.

Configuring Centralized Image Upgrades

This feature can be configured on a master controller only, and supports up to 100 simultaneous downloads. You can configure a centralized image upgrade using the WebUI or command-line interfaces.

Using the WebUI

1. Navigate to **Maintenance > Controller> Image Management**.
2. Click the **Local Configuration** tab.
3. Click the **Enable** checkbox to enable this feature. When this option is selected, the WebUI displays the following centralized image configuration parameters.

Table 162: *Centralized Image Upgrade Configuration Parameters*

Parameter	Description
Protocol	Specify the protocol used to send the software upgrade from the image server to the local controller. • TFTP • FTP • SCP
Server IP address	IP address of the image server.
Username	If you selected the FTP or SCP protocol in the Protocol field, enter the username that ArubaOS uses to connect to the image server
Password	If you selected the FTP or SCP protocol in the Protocol field, enter the password that ArubaOS uses to connect to the image server
Relative File-path	Location on the image server where the image file(s) are located
Max down-loads	Maximum number of local controllers that can simultaneously download a file from a file server. The centralized image downloading feature supports up to 100 simultaneous down-loads. If this field is left blank, ArubaOS will use its default value of 10 downloads.
Reboot auto-matically	Select this checkbox to allow the local controllers to reboot after they download their new images. **NOTE:** If you enable this option, local controllers will reboot without saving any changes to their current configuration. If you have any unsaved configuration changes on your local controller that you want to retain, do not enable this option

4. Configure the image server settings described in the table above, then click **Apply** to save your changes.
5. Click the Verify button at the bottom of the **Maintenance > Controller> Image Management > Local Configuration** page. When you verify the upgrade profile, the master controller attempts to connect to the file server, download the different images for each unique local controller and verify the validity of the image. Once controller images are “verified” by the master controller, the local controllers that are in the

upgrade target list connect to the file server, download the appropriate image, and upgrade their software to the downloaded version.

Next, specify which local controllers should download the image from the image server. You can allow all local controllers on the master to download an image from the upgrade server, or configure this feature to allow only controllers with a specified IP address or subnet to download the image. The upgrade target controllers are configured in the **Upgrade Target** section of the **Maintenance > Controller> Image Management > Local Configuration** page.

- **Allow All Targets**: To allow all local controllers associated with that master to download an image from the image server, select the **all** option in the **Upgrade Target** section.
- Select **Targets by IP address/Subnet**: To allow local controllers with a specific IP address or subnet mask to download the image:
  1. Click **New**.
  2. Enter the IP address of a controller or the subnet mask of a group of local controllers.
  3. Click **Add**.
  4. (Optional) Repeat steps 1-3 to add a new target.
  5. Click **Apply** to save your changes.

To remove a controller from the list of upgrade targets, click **Delete** by the IP address or subnet entry in the **Upgrade Targets** table. To clear the entire list of controllers in the **Upgrade Targets** table, click the **Purge the entire target list** checkbox.

In the CLI

Access the command-line interface of the master controller in config mode, and issue the following commands:

```
upgrade-profile
   auto-reboot
   filepath <filepath>
   max-downloads <1-100>
   no ...
   password  <password>
   protocol tftp|ftp|scp
   serverip  <ipaddr>
   upgrade-enable
   username <username>
```

The following commands are available in enable mode on master controllers:

```
upgrade verify
upgrade target
   all
   host <ipaddr>
   net  <subnet>
```

Viewing Controller Upgrade Statistics

The **Maintenance > Controller> Image Management > Upgrade Status** page in the WebUI and the output of the **show upgrade status** and **show upgrade configuration** commands in the command-line interface display current controller upgrade statistics.

Table 163: *All ControllersTable Data*

Column	Description
IP Address	IP address of a controller that can download images from the image file server.
Hostname	Name of the controller.
Type	Controller type (local or master)
Model	Controller model.
Version	Version of software currently running on the controller.
Upgrade Status	A controller configured to use the centralized image update feature can have one of the following upgrade status types: • **N/A**: Not applicable. Only the master controller has this status type. (Or the active master if a standby controller is configured.) • **Rebooting**: The local controller upgraded its image and is rebooting. • **Up-to-date**: The local or standby controller is running the same image as the master controller. • **Waiting, image not verified**: The local controller is waiting for the master controller to verify the images are present in the file server. • **Not Supporte**d: The local controller version is lower than ArubaOS 6.3 and does not support the upgrade feature. • **Upgraded, reboot required**: The local controller upgraded its image and a reboot is needed. A controller can have this status if the auto-reboot setting is not enabled in the upgrade profile. • **Not part of target**: The local controller image version does not match with the master and requires an upgrade, but is not part of the target upgrade list.

Managing Certificates

The controller is designed to provide secure services through the use of digital certificates. Certificates provide security when authenticating users and computers and eliminate the need for less secure password-based authentication.

There is a *default* server certificate installed in the controller to demonstrate the authentication of the controller for captive portal and WebUI management access. However, this certificate does not guarantee security in production networks. Dell*strongly* recommends that you replace the default certificate with a custom certificate issued for your site or domain by a trusted Certificate Authority (CA). This section describes how to generate a Certificate Signing Request (CSR) to submit to a CA and how to import the signed certificate received from the CA into the controller.

The controller supports client authentication using digital certificates for specific user-centric network services, such as AAA FastConnect, VPN (see Virtual Private Networks on page 337), and WebUI and SSH management access. Each service can employ different sets of client and server certificates.

During certificate-based authentication, the controller provides its server certificate to the client for authentication. After validating the controller's server certificate, the client presents its own certificate to the controller for authentication. To validate the client certificate, the controller checks the certificate revocation list (CRL) maintained by the CA that issued the client certificate. After validating the client's certificate, the controller can check the user name in the certificate with the configured authentication server (this action is optional and configurable).

When using X.509 certificates for authentication, if a banner message has been configured on the controller, it displays before the user can login. Click on a "login" button after viewing the banner message to complete the login process.

About Digital Certificates

Clients and the servers to which they connect may hold authentication certificates that validate their identities. When a client connects to a server for the first time, or the first time since its previous certificate has expired or been revoked, the server requests that the client transmit its authentication certificate. The client's certificate is then verified against the CA which issued it. Clients can also request and verify the server's authentication certificate. For some applications, such as 802.1x authentication, clients do not need to validate the server certificate for the authentication to function.

Digital certificates are issued by a CA which can be either a commercial, third-party company or a private CA controlled by your organization. The CA is trusted to authenticate the owner of the certificate before issuing a certificate. A CA-signed certificate guarantees the identity of the certificate holder. This is done by comparing the digital signature on a client or server certificate to the signature on the certificate for the CA. When CA-signed certificates are used to authenticate clients, the controller checks the validity of client certificates using certificate revocation lists (CRLs) maintained by the CA that issued the certificate.

Digital certificates employ public key infrastructure (PKI), which requires a private-public key pair. A digital certificate is associated with a private key, known only to the certificate owner, and a public key. A certificate encrypted with a private key is decrypted with its public key. For example, party A encrypts its certificate with its private key and sends it to party B. Party B decrypts the certificate with party A's public key.

Obtaining a Server Certificate

Best practices is to replace the default server certificate in the controller with a custom certificate issued for your site or domain by a trusted CA. To obtain a security certificate for the controller from a CA:

1. Generate a Certificate Signing Request (CSR) on the controller using either the WebUI or CLI.
2. Submit the CSR to a CA. Copy and paste the output of the CSR into an email and send it to the CA of your choice.
3. The CA returns a signed server certificate and the CA's certificate and public key.
4. Install the server certificate, as described in Importing Certificates on page 798.

There can be only one outstanding CSR at a time in the controller. Once you generate a CSR, you need to import the CA-signed certificate into the controller before you can generate another CSR.

In the WebUI

1. Navigate to the **Configuration > Management > Certificates > CSR** page.
2. Enter the following information:

Table 164: *CSR Parameters*

Parameter	Description	Range
`CSR Type`	Type of the CSR. You can generate a certificate signing request either with an Elliptic curve (EC) key, or with a Rivest-Shamir-Aldeman (RSA) key.	ec/rsa
`Curve name`	Length of the private/public key for ECDSA. This is applicable only if **CSR Type** is `ec`.	secp256r1/secp384r1
`Key Length`	Length of the private/public key for RSA. This is applicable only if **CSR Type** is `rsa`.	1024/2048/4096
`Common Name`	Typically, this is the host and domain name, as in www.yourcompany.com.	—
`Country`	Two-letter ISO country code for the country in which your organization is located.	
`State/Province`	State, province, region, or territory in which your organization is located.	
`City`	City in which your organization is located.	
`Organization`	Name of your organization.	
`Unit`	Optional field to distinguish a department or other unit within your organization.	
`Email Address`	Email address referenced in the CSR.	

3. Click **Generate New**.
4. Click **View Current** to display the generated CSR. Select and copy the CSR output between the BEGIN CERTIFICATE REQUEST and END CERTIFICATE REQUEST lines, paste it into an email and send it to the CA of your choice.

In the CLI

1. Run the following command:

```
crypto pki csr {rsa key_len <key_val> |{ec curve-name <key_val>} common_name <common_val>
country <country_val> state_or_province <state> city <city_val> organization <organization_
val> unit <unit_val> email <email_val>
```

2. Display the CSR output with the following command:

```
show crypto pki csr
```

3. Copy the CSR output between the BEGIN CERTIFICATE REQUEST and END CERTIFICATE REQUEST lines, paste it into an email and send it to the CA of your choice.

Obtaining a Client Certificate

You can use the CSR generated on the controller to obtain a certificate for a client. However, since there may be a large number of clients in a network, you typically obtain client certificates from a corporate CA server. For example, in a browser window, enter http://<ipaddr>/crtserv, where <ipaddr> is the IP address of the CA server.

Importing Certificates

Use the WebUI or the CLI to import certificates into the controller.



You cannot export certificates from the controller.

You can import the following types of certificates into the controller:

- Server certificate signed by a trusted CA. This includes a public and private key pair.
- CA certificate used to validate other server or client certificates. This includes only the public key for the certificate.
- Client certificate and client's public key. (The public key is used for applications such as SSH which does not support X509 certificates and requires the public key to verify an allowed certificate.)

Certificates can be in the following formats:

- X509 PEM unencrypted
- X509 PEM encrypted with a key
- DER
- PKCS7 encrypted
- PKCS12 encrypted

In the WebUI

1. Navigate to the **Configuration > Management > Certificates > Upload** page.
2. For **Certificate Name**, enter a user-defined name.
3. For **Certificate Filename**, click **Browse** to navigate to the appropriate file on your computer.
4. If the certificate is encrypted, enter the passphrase.
5. Select the **Certificate Format** from the drop-down menu.
6. Select the **Certificate Type** from the drop-down menu.
7. Click **Upload** to install the certificate in the controller.

In the CLI

Use the following command to import CSR certificates:

```
crypto pki-import {der|pem|pfx|pkcs12|pkcs7} {PublicCert|ServerCert|TrustedCA} <name>
```

The following example imports a server certificate named **cert_20** in DER format:

```
crypto pki-import der ServerCert cert_20
```

Viewing Certificate Information

In the WebUI, the Certificate Lists section of the page lists the certificates that are currently installed in the controller. Click **View** to display the contents of a certificate.

To view the contents of a certificate with the CLI, use the following commands:

Table 165: *Certificate Show Commands*

Command	Description
`show crypto-local pki trustedCAs [<name>]< [attribute>]`	Displays the contents of a trusted CA certificate. If a name is not specified, all CA certificates imported into the controller are displayed. If name and attribute are specified, then only the attribute in the certificate are displayed. Attributes can be CN, validity, serial-number, issuer, subject, public-key.
`show crypto-local pki serverCerts [<name>] [<attribute>]`	Displays the contents of a server certificate. If a name is not specified, all server certificates imported into the controller are displayed.
`show crypto-local pki publiccert [<name>] [<attribute>]`	Displays the contents of a public certificate. If a name is not specified, all public certificates imported into the controller are displayed.

Imported Certificate Locations

Imported certificates and keys are stored in the following locations in flash on the controller:

Table 166: *Imported Certificate Locations*

Location	Description
/flash/certmgr/trustedCAs	Trusted CA certificates, either for root or intermediate CAs. Best practices is to import the certificate for an intermediate CA, you also import the certificate for the signing CA.
/flash/certmgr/serverCerts	Server certificates. These certificates must contain both a public and private key (the public and private key must match). You can import certificates in PKCS12 and X509 PEM formats, but they are stored in X509 PEM DES encrypted format.
/flash/certmgr/CSR	Temporary certificate signing requests (CSRs) that have been generated on the controller and are awaiting a CA to sign them.
/flash/certmgr/publiccert	Public key of certificates. This allows a service on the controller to identify a certificate as an allowed certificate.

Checking CRLs

A CA maintains a CRL that contains a list of certificates that have been revoked before their expiration date. Expired client certificates are not accepted for any user-centric network service. Certificates may be revoked because certificate key has been compromised or the user specified in the certificate is no longer authorized to use the key.

When a client certificate is being authenticated for a user-centric network service, the controller checks with the appropriate CA to make sure that the certificate has not been revoked.



The controller does not support download of CRLs.

Certificate Expiration Alert

The certificate expiration alert sends alerts when installed certificates, which correspond to trust chains, OCSP responder certificates, and any other certificates installed on the device. By default, the system sends this alert 60 days before the expiration of the installed credentials. This alert is then repeated periodically on a weekly or biweekly basis. This alerts consist of two SNMP traps:

- wlsxCertExpiringSoon
- wlsxCertExpired

Chained Certificates on the RAP

Chained certificates on the RAP (that is, certificates from a multi-level PKI) need to be in a particular order inside the file. The RAP's certificate must be first, followed by the certificate chain in order, and then followed by the private key for the certificate. For example, with a root CA, a single intermediate CA, and a root CA, the PEM or PKCS12 file must contain the following parts, in this order:

1. RAP Certificate
2. Intermediate CA
3. Root CA
4. Private key



If this order is not followed, certificate validation errors occur. This order also applies to server certificates.

Support for Certificates on USB Flash Drives

This release now supports the USB storing of the RAP certificate. This ensures that the RAP certificate is activated only when the USB with the corresponding certificate is connected to the RAP. Likewise, the RAP certificate is deactivated when the USB is removed from the RAP. In this case, the USB that is connected to the RAP is an actual storage device and does not act as a 3G/4G RAP.

The RAP supports only PKCS12-encoded certificates that are present in the USB. This certificate contains all the information that is required for creating the tunnel including the private key, RAP certificate with the chain of certificates and the trusted CA certificate. There is a limit of three supported intermediate CAs and the common name for the RAP certificate must be the MAC address of the RAP in the colon format.



If you have an activated RAP that is using USB storage for the certificate, and you remove the USB storage, the RAP drops the tunnel. This is by design. However, for the RAP to re-establish the tunnel it has to be power cycled. It does not matter if you reinsert the USB storage before or after the power cycle as long as you power cycle it.

Marking the USB Device Connected as a Storage Device

If the AP provisioning parameter "usb-type" contains the value "storage," this indicates that the RAP will retrieve certificates from the connected USB flash drive.

RAP Configuration Requirements

The RAP needs to have one additional provisioning parameter, the pkcs12_passphrase, which can be left untouched or can store an ACSII string. The string assigned to this parameter is used as the passphrase for decoding the private key stored.



If you have an activated RAP that is using USB storage for the certificate, and you remove the USB storage, the RAP drops the tunnel. This is by design. However, for the RAP to re-establish the tunnel it has to be power cycled. It does not matter if you reinsert the USB storage before or after the power cycle as long as you power cycle it.

When the RAP successfully extracts all the information including the CA certificate, the RAP certificate and the RAP private key using the passphrase from the provisioning parameter, it successfully establishes the tunnel.

Configuring SNMP

Dell controllers support versions 1, 2c, and 3 of Simple Network Management Protocol (SNMP) for reporting purposes only. In other words, SNMP cannot be used for setting values in a Dell system in the current ArubaOS version. *MIB Reference Guide* for information about the Dell MIBs and SNMP traps..



Dell-specific management information bases (MIBs) describe the objects that can be managed using SNMP. See the *Dell Networking W-Series ArubaOS MIB Reference Guide* for information about the Dell MIBS and SNMP traps.

SNMP Parameters for the Controller

You can configure the following SNMP parameters for the controller.

Table 167: *SNMP Parameters for the Controller*

Field	Description
Host Name	Host name of the controller.
System Contact	Name of the person who acts as the System Contact or administrator for the controller.
System Location	String to describe the location of the controller.
Read Community Strings	Community strings used to authenticate requests for SNMP versions before version 3. **NOTE:** This is needed only if using SNMP v2c and is not needed if using version 3.
Enable Trap Generation	Enables generation of SNMP traps to configured SNMP trap receivers. Refer to the list of traps in the “SNMP traps” section below for a list of traps that are generated by the controller.
Trap receivers	Host information about a trap receiver. This host needs to be running a trap receiver to receive and interpret the traps sent by the Dell controller. Configure the following for each host/trap receiver: • IP address • SNMP version: can be 1, 2c, or 3. • Type: Trap or Inform (SNMPv2c or SNMPv3 only) • Engine ID: (SNMPv3 only) • Security string • UDP port on which the trap receiver is listening for traps. The default is the UDP port number 162. This is optional, and will use the default port number if not modified by the user.
If you are using SNMPv3 to obtain values from the controller, you can configure the following parameters:	
User name	A string representing the name of the user.
Authentication protocol	An indication of whether messages sent on behalf of this user can be authenticated, and if so, the type of authentication protocol used. This can take one of the two values: • MD5: HMAC-MD5-96 Digest Authentication Protocol • SHA: HMAC-SHA-96 Digest Authentication Protocol

Field	Description
Authentication protocol password	If messages sent on behalf of this user can be authenticated, the (private) authentication key for use with the authentication protocol. This is a string password for MD5 or SHA depending on the choice above.
Privacy protocol	An indication of whether messages sent on behalf of this user can be protected from disclosure, and if so, the type of privacy protocol which is used. This takes the value DES (CBC-DES Symmetric Encryption Protocol).
Privacy protocol password	If messages sent on behalf of this user can be encrypted/decrypted with DES, the (private) privacy key for use with the privacy protocol.

Follow the steps below to configure a controller's basic SNMP parameters.

In the WebUI

1. Navigate to the **Configuration > Management > SNMP** page.
2. If the controller will be sending SNMP traps, click **Add** in the Trap Receivers section to add a trap receiver.
3. If you are using SNMPv3 to obtain values from the controller, click **Add** in the SNMPv3 Users section to add a new SNMPv3 user.
4. Click **Apply**.

In the CLI

```
hostname name
syscontact name
syslocation string
snmp-server community string
snmp-server enable trap
snmp-server engine-id engine-id
snmp-server host ipaddr version {1|2c|3} string [udp-port number]
snmp-server trap source ipaddr
snmp-server user name [auth-prot {md5|sha} password priv-prot DES password
```



Earlier versions of ArubaOS supported SNMP on individual APs. This feature is not supported by this version of ArubaOS.

Enabling Capacity Alerts

Use the capacity alert feature to set controller capacity thresholds which, when exceeded, will trigger alerts. The controller will send a *wlsxThresholdExceeded* SNMP trap and a syslog error message when the controller has exceeded a set percentage of the total capacity for that resource. A *wlsxThresholdCleared* SNMP trap and error message will be triggered if the resource usage drops below the threshold once again.

The following table describes the thresholds that can be configured with this feature.

Table 168: *Capacity Alert Thresholds*

Threshold	Description
controlpath-cpu	Set an alert threshold for controlpath CPU capacity. The <percentage> parameter is the percentage of the total controlpath CPU capacity that must be exceeded before the alert is sent. The default threshold for this parameter is 80%.
controlpath-memory	Set an alert threshold for controlpath memory consumption. The <percentage> parameter is the percentage of the total memory capacity that must be exceeded before the alert is sent. The default threshold for this parameter is 80%.
datapath-cpu	Set an alert threshold for datapath CPU capacity. The <percentage> parameter is the percentage of the total datapath CPU capacity that must be exceeded before the alert is sent. The default threshold for this parameter is 30%.
no-of-APs	The maximum number of APs that can be connected to a controller is determined by that controller's model type and installed licenses. Use this command to trigger an alert when the number of APs currently connected to the controller exceeds a specific percentage of its total AP capacity. The default threshold for this parameter is 80%.
no-of-locals	Set an alert threshold for the master controller's capacity to support remote nodes and local controllers. A master controller can support a combined total of 256 remote nodes and local controllers. The <percentage> parameter is the percentage of the total master controller capacity that must be exceeded before the alert is sent. The default threshold for this parameter is 80%.
total-tunnel-capacity	Set an alert threshold for the controller's tunnel capacity. The <percentage> parameter is the percentage of the controller's total tunnel capacity that must be exceeded before the alert is sent. The default threshold for this parameter is 80%
user-capacity	Set an alert threshold for the controller's user capacity. The <percentage> parameter is the percentage of the total resource capacity that must be exceeded before the alert is sent. The default threshold for this parameter is 80%.

In the WebUI

1. Navigate to the **Configuration > Management > Threshold**page.
2. Modify the capacity percentages for any of the thresholds described in "Capacity Alert Thresholds" on page 657.
3. Click**Apply** to save your settings.

In the CLI

4. To configure this feature, access the command-line interface in config mode and issue the following commands:

```
threshold
   controlpath-cpu <percentage>
   controlpath-memory <percentage>
   datapath-cpu <percentage>
   no-of-APs <percentage>
   no-of-locals <percentage>
   total-tunnel-capacity <percentage>
   user-capacity <percentage>
```

Examples

The following command configures a new alert threshold for controlpath memory consumption:

```
(host) (config) #threshold datapath-cpu 90
```

If this threshold is exceeded then subsequently drops below the 90% threshold, the controller would send the following two syslog error messages.

```
May 14 13:13:58  nanny[1393]: <399816> <ERRS> |nanny|  Resource 'Control-Path Memory' has gone
above  90%  threshold, value : 93
May 14 13:16:58  nanny[1393]: <399816> <ERRS> |nanny|  Resource 'Control-Path Memory' has come
below  90%  threshold, value : 87
```

Configuring Logging

This section outlines the steps required to configure logging on a controller.

For each category or subcategory of message, you can set the logging level or severity level of the messages to be logged. Table 169 summarizes these categories:

Table 169: *Software Modules*

Category/Subcategory	Description
Network	Network messages
all	All network messages
packet-dump	Protocol packet dump messages
mobility	Mobility messages
dhcp	DHCP messages
System	System messages
all	All system messages
configuration	Configuration messages
messages	Messages
snmp	SNMP messages
webserver	Web server messages
security	Security messages
all	All security messages
aaa	AAA messages
firewall	Firewall messages
packet-trace	Packet trace messages
mobility	Mobility messages
vpn	VPN messages

Category/Subcategory	Description
dot1x	802.1x messages
ike	IKE messages
webserver	Web server messages
Wireless	Wireless messages
all	All wireless messages
User	User messages
all	All user messages
captive-portal	Captive portal user messages
vpn	VPN messages
dot1x	802.1x messages
radius	RADIUS user messages

For each category or subcategory, you can configure a logging level. Table 170 describes the logging levels in order of severity, from most to least severe.

Table 170: *Logging Levels*

Logging Level	Description
Emergency	Panic conditions that occur when the system becomes unusable.
Alert	Any condition requiring immediate attention and correction.
Critical	Any critical conditions such as a hard drive error.
Errors	Error conditions.
Warning	Warning messages.
Notice	Significant events of a non-critical and normal nature.
Informational	Messages of general interest to system users.
Debug	Messages containing information useful for debugging.

The default logging level for all categories is Warning. You can also configure IP address of a syslog server to which the controller can direct these logs.

In the WebUI

1. Navigate to the **Configuration > Management > Logging > Servers** page.
2. To add a logging server, click **New** in the Logging Servers section.
3. Click **Add** to add the logging server to the list of logging servers. Ensure that the syslog server is enabled and configured on this host. Click **Apply**.
4. To select the types of messages you want to log, select the **Levels** tab.

5. Select the category or subcategory to be logged.
6. To select the severity level for the category or subcategory, scroll to the bottom of the page. Select the level from the Logging Level drop-down menu. Click **Done**.
7. Click **Apply** to apply the configuration.

In the CLI

```
logging <ipaddr>
logging level <level> <category> [subcat <subcategory>]
```

Syslog operates over UDP and is connectionless. Therefore, it is not possible for the controller to recognize a failure of the syslog server or the network path to the syslog server. By establishing an IPsec tunnel between the controller and the syslog server, (see Planning a VPN Configuration) it is possible to indirectly track the status of the syslog server link.

After a failure occurs, the network administrator has to manually re-synchronize log files by copying them from the controller to the syslog server. Use the **tar logs** CLI command to create an archive of all local logs, then use the **copy** CLI command to copy this archive to an external server. Log space is limited on the controller, and depending on how long the outage lasted some local logs may be overwritten.

Enabling Guest Provisioning

The Guest Provisioning feature lets you manage guests who need access to your company's wireless network. This section describes how to:

- Design and configure the Guest Provisioning page – Using the WebUI, the network administrator designs and configures the Guest Provisioning page that is used to create a guest account.
- Configure a guest provisioning user – The network administrator configures one or more guest provisioning users. A guest provisioning user, such as a front desk receptionist, signs in guests at your company.
- Using the Guest Provisioning page – The Guest Provisioning page is used by the guest provisioning user to create guest accounts for people who are visiting your company.

Configuring the Guest Provisioning Page

Use the Guest Provisioning Configuration page to create the Guest Provisioning page. This configuration page consists of three tabs: Guest Fields, Page Design and Email. You configure the information on all three tabs to create a Guest Provisioning page.

- Guest Fields tab—lets you select the fields that appear on the Guest Provisioning page.
- Page Design tab—lets you specify the company banner, heading, and text and background colors that appear on the Guest Provisioning page.
- Email tab—lets you specify an email to be sent to the guest or sponsor (or both). Email messages can be sent automatically at account creation time and also may be sent manually by the administrator from the Guest Provisioning page.

In the WebUI



You can only create and design the Guest Provisioning page in the WebUI.

This section describes how to design a Guest Provisioning page using all three tabs.

Configuring the Guest Fields

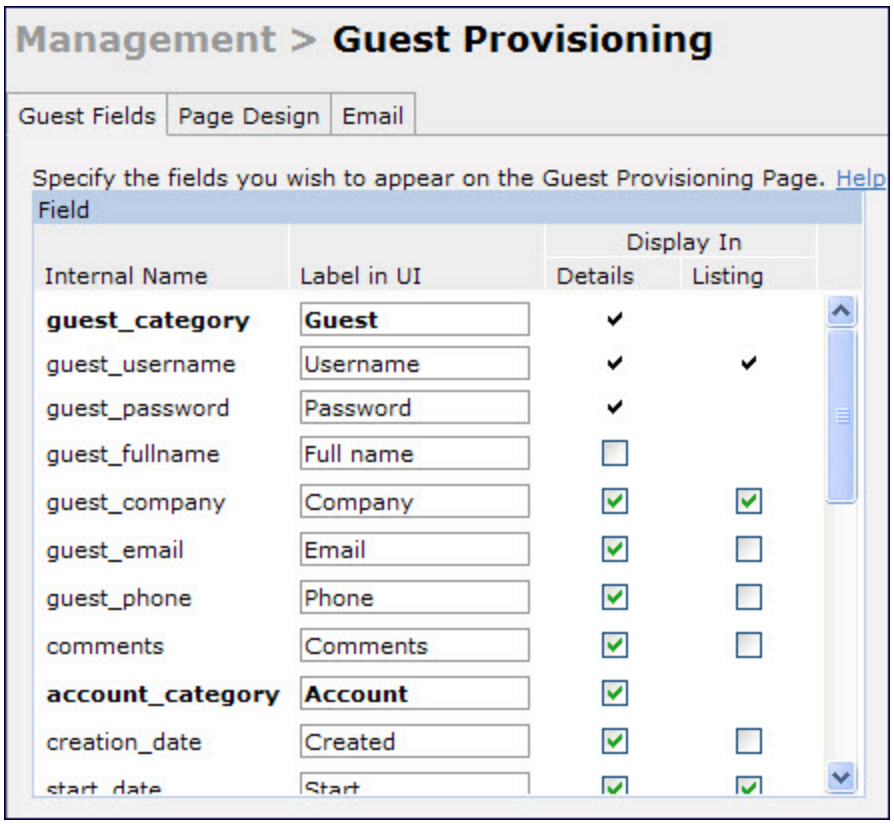
1. Navigate to the **Configuration > Management > Guest Provisioning** page. The Guest Provisioning configuration page displays with the Guest Fields tab on top. This tab contains the following columns:
   - Internal Name—The unique identifier that is mapped to the label in the UI.
   - Label in UI—A customizable string that displays in both the main listing pane and details sheet on the Guest Provisioning page.
   - Display in Details—Fields with selected checkboxes appear in the Show Details popup-window.

NOTE

If the guest_category, account_category, sponsor_category and optional_category fields are not checked, their respective sections do not appear on the Guest Provisioning page.

   - Display in Listing—Fields with selected checkboxes appear as columns in the management user summary page.

Figure 164 *Guest Provisioning Configuration Page—Guest Fields Tab*



2. Select the checkbox next to each field, described in Table 171, that you want to appear on the Guest Provisioning page. Optionally, you can customize the label that displays in the UI.
3. Click **Preview Current Settings** to view what the Guest Provisioning page looks like while you are designing it.
4. To save changes, click **Apply**.

NOTE

Best practices is to check the **Display in Listing** field for only the most essential fields, so that the Guest Provisioning user does not have to scroll the guest listing horizontally to see all the columns.

Table 171: *Guest Provisioning—Guest Field Descriptions*

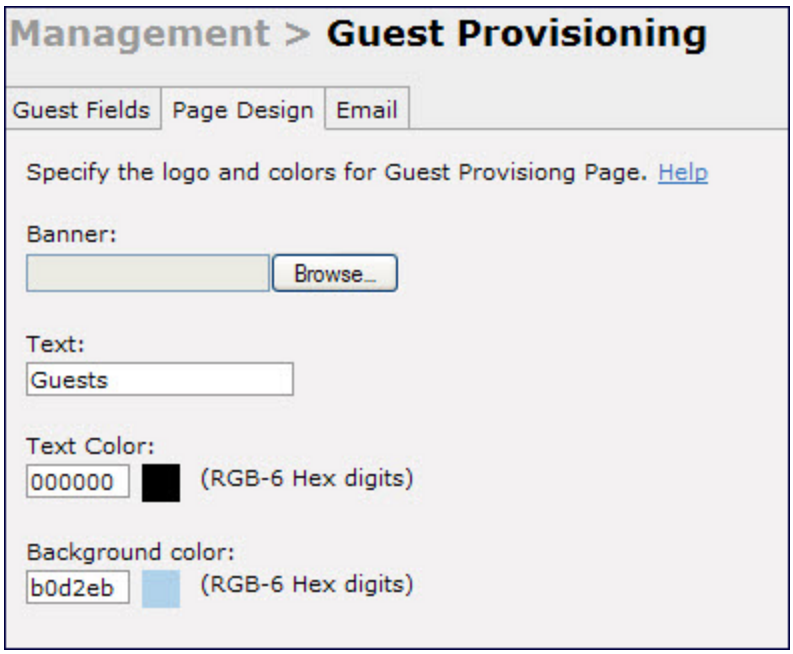
Guest Field	Description
guest_category	A guest is the person who needs guest access to the company's wireless network. This is the label on the Guest Provisioning page for the guest information.
guest_username	Username for the guest.
guest_password	Password for the guest. (Must contain at least 1-6 characters and at least one digit.)
guest_fullname	Full name of the guest.
guest_company	Name of the guest's company.
guest_email	Guest's Email address.
guest_phone	Guest's phone number
comments	Optional comments about the guest's account status, meeting schedule and so on.
account_category	This is the label on the Guest Provisioning page for the account information.
creation-date	Date the account is created.
start_date	Date the guest account begins.
end_date	Date the guest account ends.
grantor	The username of the person of who created the guest account.
grantor_role	The authentication role of the grantor.
sponsor_category	A sponsor is the guest's primary contact for the visit. This is the label in the Guest Provisioning page for the sponsor information.
sponsor_username	
	Sponsor's work department
sponsor_email	Sponsor's Email address.
optional_category	This is the label in the Guest Provisioning page for the information in the optional fields that follow. **NOTE:** The optional_category field can be used for another person, for example a "Supervisor." You can enter username, full name, department and Email information into the optional fields. Or, you can use this category for some other purpose.
optional_field_1	optional_field_1 description
optional_field_2	optional_field_2 description
optional_field_3	optional_field_2 description
optional_field_4	optional_field_2 description

Configuring the Page Design

The Page Design tab lets you specify the company banner, heading, and text and background colors that appear on the Guest Provisioning page.

1. Navigate to the **Configuration > Management > Guest Provisioning** page and select the **Page Design** tab.

Figure 165 *Guest Provisioning Configuration Page—Page Design Tab*



2. Enter the filename which contains the company banner in the **Banner** field. Or, click **Browse** to search for the filename



Best practices is to use a logo or banner image that is 600 x 100 pixels (width x height). The WebUI does not apply the size restrictions when you upload an image file, but the image is resized to 600 x 100 pixels when it displays or is printed.

3. Enter the label for the guest listing (the one you used in the Guest Fields tab) in the **Text** field.
4. Enter the hex value for the color of the text in the **Text Color** field. The text in the header of the guest listing displays in this color.
5. Enter the hex value for the color of the background in the **Background color** field. This determines the color of the header of the guest listing.
6. Click **Preview Current Settings** to preview the Guest Provisioning page while you are designing it.
7. To save changes, click **Apply**.

Configuring Email Messages

You can specify an email to be sent to the guest or sponsor (or both). Email messages can be sent automatically at account creation time or sent manually by the network administrator or guest provisioning user from the Guest Provisioning page at any time.

1. Specify the SMTP server and port that processes the guest provisioning (also known as guest access) email. You can complete this step using the WebUI or CLI commands:
   - Configuring the SMTP Server and Port in the WebUI on page 810
   - Configuring an SMTP server and port in the CLI on page 810

2. Create the email messages. Complete this step using the WebUI:
   Creating Email Messages in the WebUI on page 810

Configuring the SMTP Server and Port in the WebUI

1. Navigate to the **Configuration > Management > SMTP**page**.**
2. Enter the IP address of the SMTP server to which the controller sends the guest provisioning email in the **IP Address of SMTP** server field.
3. Enter the number of the port through which the guest provisioning email passes in the **Port** field.
4. Click **Apply** and then **Save Configuration**.

Configuring an SMTP server and port in the CLI

The following command creates a guest-access email and sends guest user email through SMTP server IP address 1.1.1.1 on port 25.

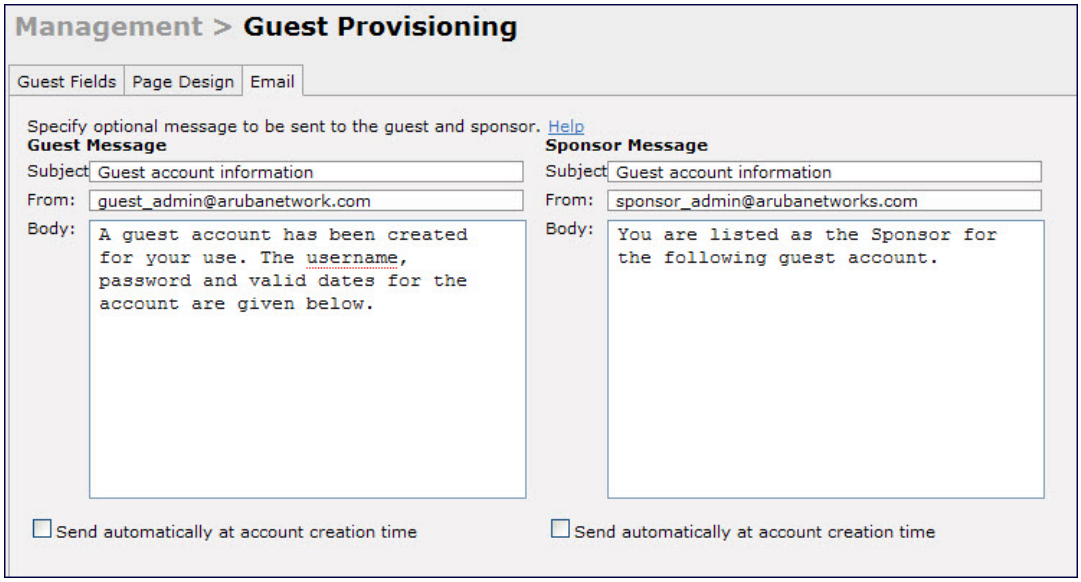
```
(host) (config) #guest-access-email
(host) (Guest-access Email) #
(host) (Guest-access Email) #smtp-port 25
(host) (Guest-access Email) #smtp-server 1.1.1.1
```

Creating Email Messages in the WebUI

After you configured the SMTP server and port, follow these steps:

1. Navigate to the **Configuration > Management > Guest Provisioning** page and select the **Email** tab.

Figure 166 *Guest Provisioning Configuration Page—Email Tab*



2. To create a message for a guest or sponsor, customize the text in the **Subject**, **From**,s and **Body** fields as needed for both the Guest message and Sponsor message.
3. Optionally, select the **Send automatically at account creation time** checkbox when you want an email message to be sent to the guest and/or sponsor alerting them that a guest account has just been created.



Regardless of whether you select this option, the person responsible for managing the Guest Provisioning page may choose to send this email message manually at any time.

Figure 167 shows a sample email message that is sent to the guest after the guest account is created.

Figure 167 *Sample Guest Account Email – Sent to Sponsor*

```
Sent: Monday, February 09, 2009 12:59 PM
To: John Smith
Subject: Guest account information

A guest account has been created for your use. The username, password and
valid dates for the account are given below.
================================================
Username:  guest3518444
Password:  hqtehjc1936850
Guest Name:
Guest Company:  MyCompany
Guest Email:  JSmith@MyCompany.com
Guest Phone:
Sponsor Email:  DJones@AcmeCompany.com
Start Date:  Mon Feb  9 18:46:00 2009
Expiration Date:  Mon Feb  9 19:46:00 2009
```

4. To save changes, click **Apply**.

Configuring a Guest Provisioning User

The guest provisioning user has access to the Guest Provisioning Page (GPP) to create guest accounts within your company. The guest provisioning user is usually a person at the front desk who greets guests and creates guest accounts. Depending upon your needs, there are three ways to configure and authenticate a guest provisioning user:

- Username and Password authentication — Allows you to configure a user in a guest provisioning role.
- Smart Card authentication
  - Static authentication —Uses a configured certificate name and serial number to derive the user role. This authentication process uses a previously configured certificate name and serial number to derive the user role. This method does not use and external authentication server.
  - Authentication server — Uses an external authentication server to derive the management role. This is helpful if there is a large number of users who need to be deployed as guest provisioning users.

You can use the WebUI or CLI to create a Guest Provisioning user.

In the WebUI

This section describes how to configure a guest provisioning user. All three methods are described.

Username and Password Authentication Method

1. Navigate to the **Configuration > Management > Administration** page.
2. In the Management Users section, click **Add**.
3. In the Add User page select **Conventional User Accounts**.
4. In the **User Name** field, enter the name of the user who you want to configure as a guest provisioning user.
5. In the **Password** and **Confirm Password** fields, enter the user's password and reconfirm it.
6. From the **Role** drop-down menu, select **guest-provisioning**.
7. Click **Apply**.

Static Authentication Method



Before using this method, make sure that the correct CA certificate is uploaded to the controller.

1. Navigate to the **Configuration > Management > Administration** page.
2. In the Management Users section, click **Add**.
3. In the **Add User** page, select **Certificate Management**.
4. Make sure that the **Use external authentication server to authenticate** check box is unchecked.
5. In the **Username** field, enter the name of the user who you want to configure as a guest provisioning user.
6. In the **Role**field, select **guest-provisioning** from the drop-down list.
7. Enter client certificate serial number in the **Client Certificate Serial No.** field.
8. Select the CA certificate you want to use from the **Trusted CA Certificate Name** drop-down menu.
9. Click **Apply**.

Smart Card Authentication Method

1. Navigate to the **Configuration > Management > General** page.
2. In the **WebUI Management Authentication Method** section, select **Client Certificate**.
3. Click **Apply**.
4. Navigate to the **Configuration > Management > Administration** page.
5. In the **Management Authentication Servers** section, select **guest-provisioning** from the **Default Role** drop-down menu.
6. Select the **Mode** checkbox.
7. Select the server group from the **Server Group** drop-down menu.
8. Click **Apply**.
9. In the **Management Users** section, click **Add** to display the **Configuration > Management > Add User** page.
10. Select **Certificate Management**, **WebUI Certificate** and **Use external authentication server to authenticate.**
11. Select the trusted CA certificate you want to use from the **Trusted CA Certificated Name** drop-down menu.
12. Click **Apply** and **Save Configuration**.

In the CLI

Username and Password Method

This example creates a user named Paula and assigns her the role of guest provisioning.

```
(host) (config)# mgmt-user Paula guest-provisioning
```

Static Authentication Method

This example uses the CA certificate **mycertificate** with the serial number 1234 to authenticate user Laura in the guest provisioning role.

```
(host) (config)# mgmt-user webui-cacert mycertificate serial 1234 Laura guest-provisioning
```

Smart Card Authentication Method

This example shows that using previously configured certificate (1234), authentication and authorization are automatically configured using an authentication server.

```
(host) (config) #web-server mgmt-auth username/password certificate
(host) (config) #mgmt-user webui-cacert <certificate_name>
(host) (config) #aaa authentication mgmt
(host) (config) # server-group "internal"
(host) (config) #mgmt-user webui-cacert default
(host) (config) #mgmt-user webui-cacert 1234
```

Customizing the Guest Access Pass

In the WebUI, you can customize the pop-up window that displays the guest account information. You may want to do this before the Guest Provisioning user creates guest accounts.

1. Navigate to the **Configuration > Security > Access Control > Guest Access** page.
2. Click **Browse** to insert a logo or other banner information on the window.

NOTE

Best practices is to use a logo or banner image that is 600 x 100 pixels (width x height). The WebUI does not apply the size restrictions when you upload an image file, but the image is resized to 600 x 100 pixels when it displays or is printed.

3. You can enter text for the Terms and Conditions portion of the window.
4. Click **Submit** to save your changes. Click **Preview Pass** to preview the window. (See Figure 168.)

Figure 168 *Customized Guest Account Information Window*



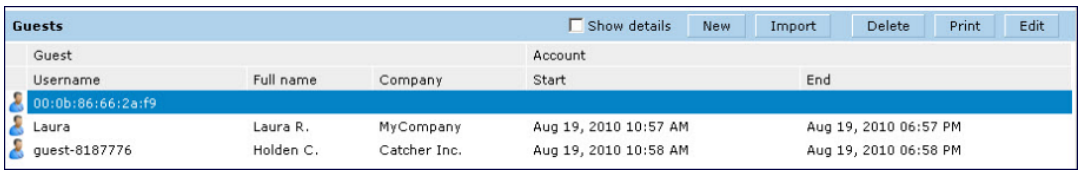
Creating Guest Accounts

After the Guest Provisioning user is created, that person can log in to the controller using the preconfigured username and password. The Guest Provisioning page displays. (See .) This is a sample page as the fields may differ based on how the network administrator designed the page.

NOTE

Starting with ArubaOS 3.4 release, a guest user account that is created by a guest provisioning user can only be viewed, modified or deleted by the guest provisioning user who created the account or the network administrator. A guest user account that is created by the network administrator can only be viewed, modified or deleted by the network administrator.

Figure 169 *Creating a Guest Account—Guest Provisioning Page*





If you do not want multiple guest users to share the same guest account concurrently, navigate to the Captive Portal Authentication and select the “Allow only one active user session” option. If a guest user authenticates successfully but the controller detects there is already a guest session with the same guest username, the second login is rejected.

Guest Provisioning User Tasks

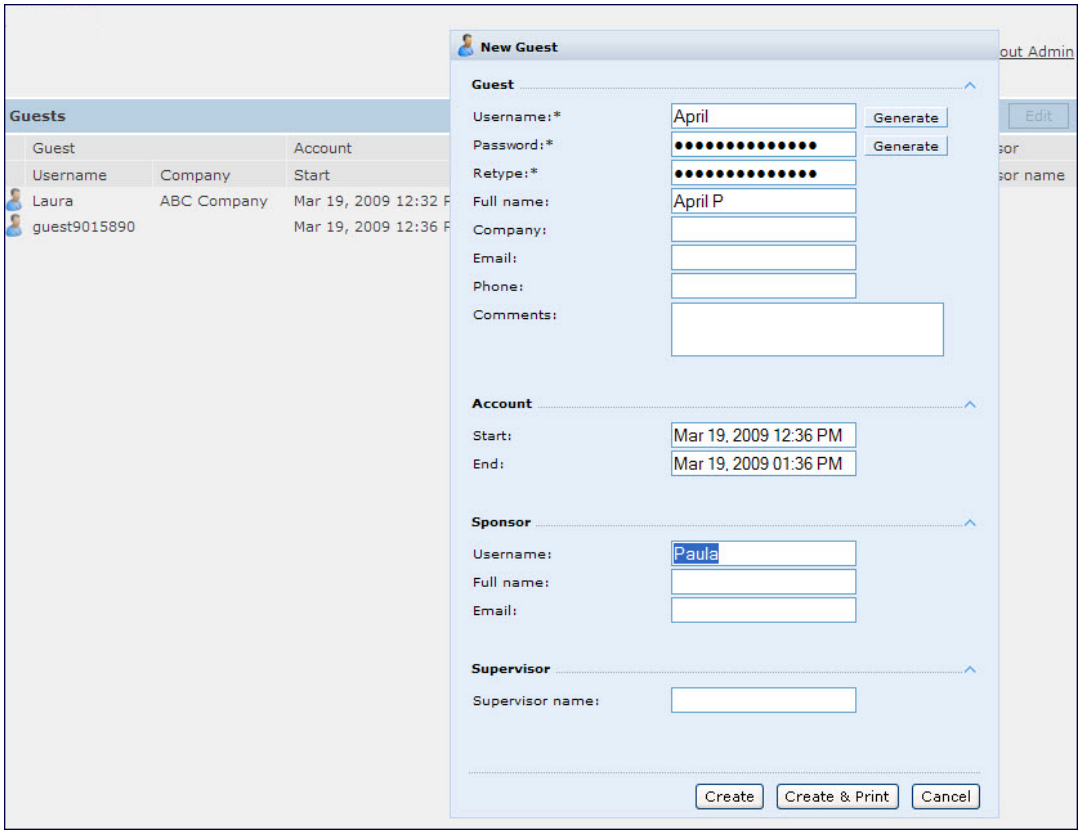
The Guest Provisioning user creates guest accounts by filling in information on the Guest Provisioning page. Tasks include creating, editing, manually sending email, enabling, printing, disabling and deleting guest accounts. The Guest Provisioning user can also manually send emails to either the guest or sponsor.

To create a new guest account, the Guest Provisioning user clicks **New** to display the New Guest window. (See Figure 170.) After filling in information into the fields, click **Create**. The guest account now displays on the Guest Provisioning page.

If you manually configure the user name and password, note the following:

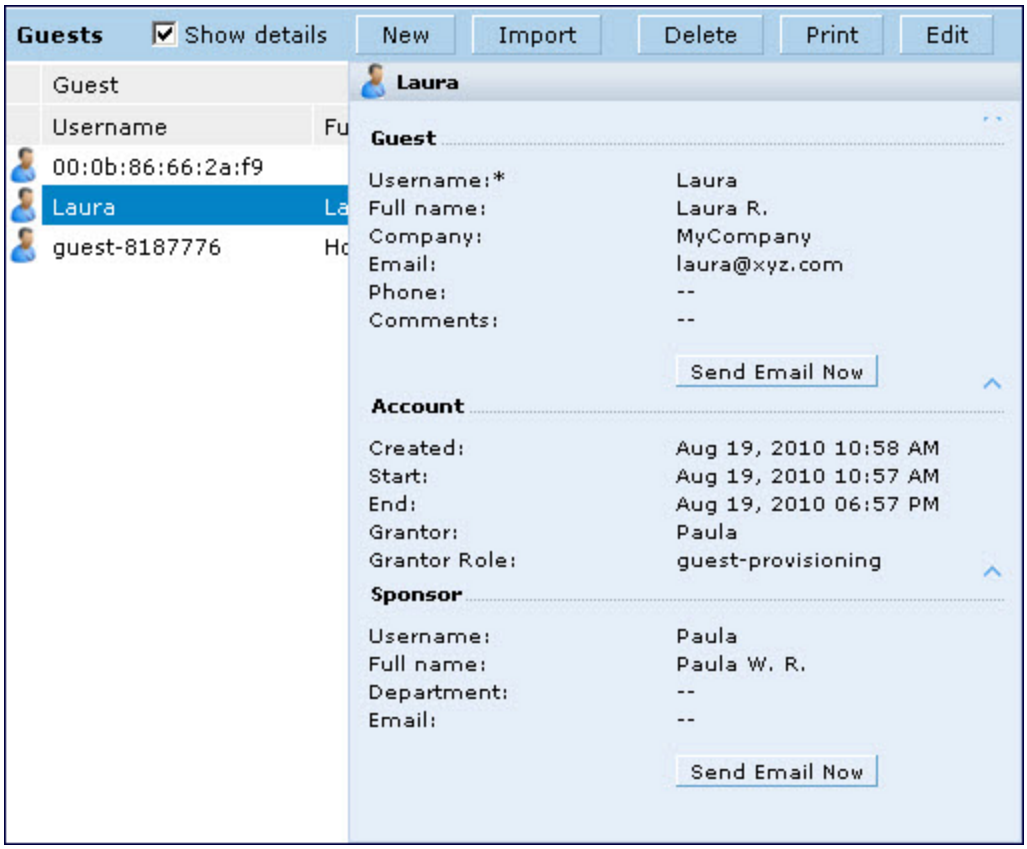
- User name entries support alphanumeric characters, however the percent sign (%) and trailing the back slash are not allowed.
- Passwords must have a minimum of six characters. You can use special characters for the password.
- Click on the **Account Start** and **End** fields to change the account start and end times. The default account start to end time setting is eight hours.

Figure 170 *Creating a Guest Account—New Guest Window*



To see details about an existing user account, highlight an existing account and select the **Show Details** checkbox. The Show Details popup-window displays. (See Figure 171.) The Guest Provisioning user can send out Email from this window to either the guest or the sponsor. When you send an email from the Details pop-up window, a pop-up message confirming that the email was successfully processed displays

Figure 171 *Creating a Guest Account—Show Details Pop-up Window*



Importing Multiple Guest Entries

The Guest Provisioning user can manually create individual guest entries, as previously described, or import multiple guest entries into the database from a CSV file. This is useful and more efficient if you want to enter multiple guest entries at once. To import multiple guest entries, you need to:

1. Create a CSV file that contains the guest entries
2. Import the CSV file into the database

Creating Multiple Guest Entries in a CSV File

Create a CSV file that contains multiple guest entries. Each field in an entry needs to be separated by a comma and each entry needs to end with a carriage return. The order of the fields is:

- Guest's first name (required)
- Guest's last name (required)
- Guest's email address (optional)
- Guest's phone number (optional)
- Guest's user ID (optional)
- Guest's password (optional)
- Sponsor's first name (optional)
- Sponsor's last name (optional)
- Sponsor's email address (optional)

See Figure 172 for an example of how guest entries need to be formatted in a CSV file.

Figure 172 *CVS File Format—Guest Entries Information*

```
Gene,Phineas,gphineas@arubanetworks.com,(415)555-1212,guest-
gwang,abcdefg,Jane,Smith,jsmith@arubanetworks.com¶
Caulfield,Holden¶
John,Galt,,,guestlll0¶
```

Note the following limitations when creating guest entries in a CVS file:

- None of the field values can have a comma
- There is no format checking on field. Only the **`local-userdb-guest`** CLI command will validate proper format.
- Any extra columns, beyond the 9th column, are discarded.
- The WebUI only supports characters that the CLI supports.
- If a guest's user ID is not provided, then it is automatically generated based on the numeric suffix in the Import Guest List window. See Figure 173.
- We recommend a maximum of 250 entries per CSV file.

Importing the CSV File into the Database

To import a CSV file that contains multiple guest entries, the Guest Provisioning user must follow these steps:

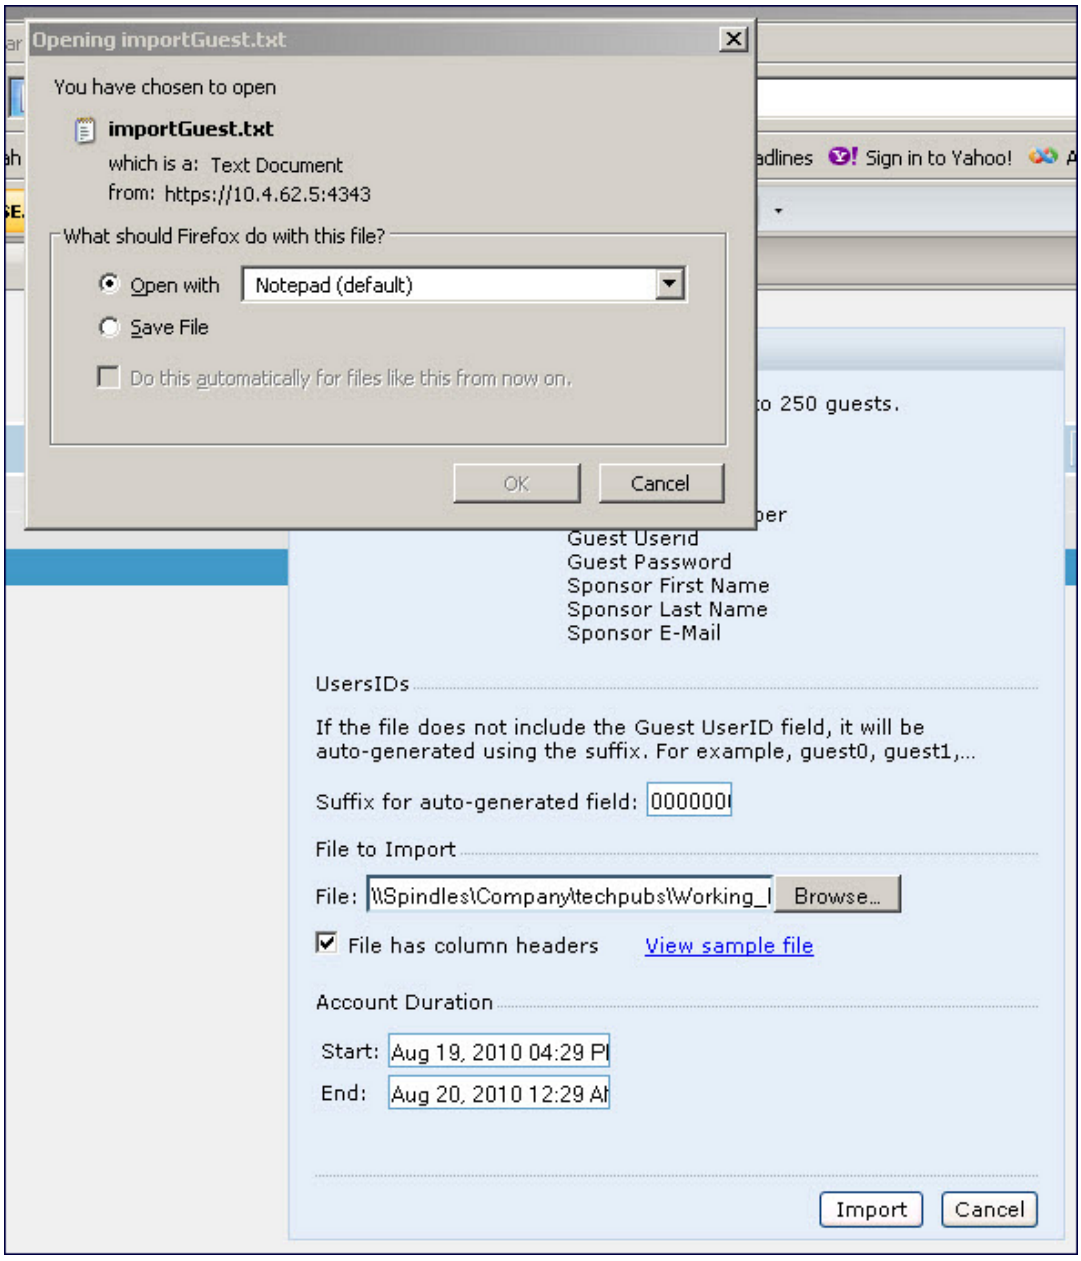
1. Log in to the WebUI using the username and password assigned to the Guest Provisioning user.
2. Click on **Import**. The **Import Guest List** pop-up window displays. See Figure 173.

Figure 173 *Importing a CSV file that contains Guest Entries*



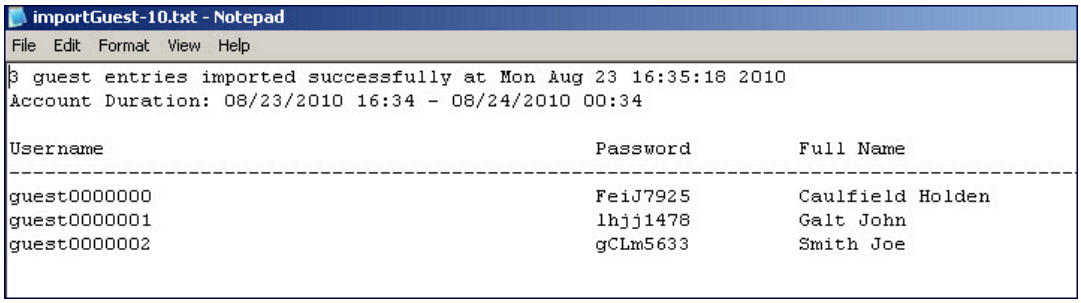
3. Click **Browse** to locate for the CSV file you want to import.
4. Click **Import**. A window displays that lets you open CSV file in text format. (See Figure 174.) Open the text file to see a summary of the number of users and error messages if users are not imported.

Figure 174 *Displaying the Guest Entries Log File*



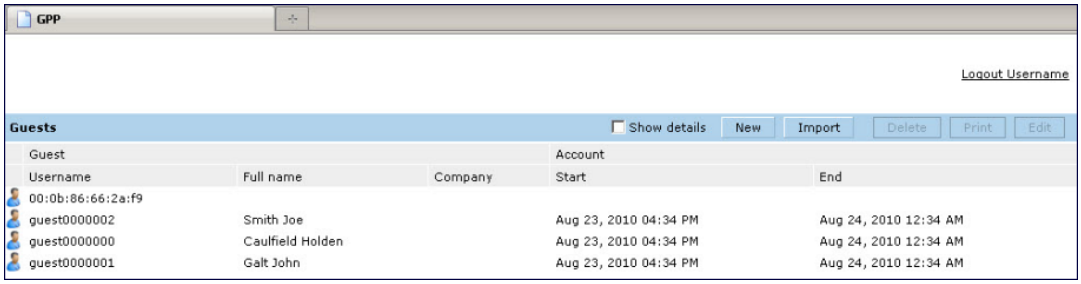
5. Click **Import**. A window displays that lets you open CSV file in text format. (See Figure 174.)
6. Open the text file. (See Figure 175.) Note that because no user ID is entered in the CSV file, a guest ID (username) is automatically generated based on the default value in the **Suffix for auto-generated** field. Make changes or corrections to the guest entry information in text file. A user can also change the start time and end time from this window. Save and exit the file.

Figure 175 *Viewing and Editing Guest Entries in the Log File*



7. Click **Cancel** to close the **Import Guest List** window. Guest entries are now displayed in the Guest Provisioning page.

Figure 176 *Viewing Multiple Imported Guest Entries—Guest Provisioning Page*

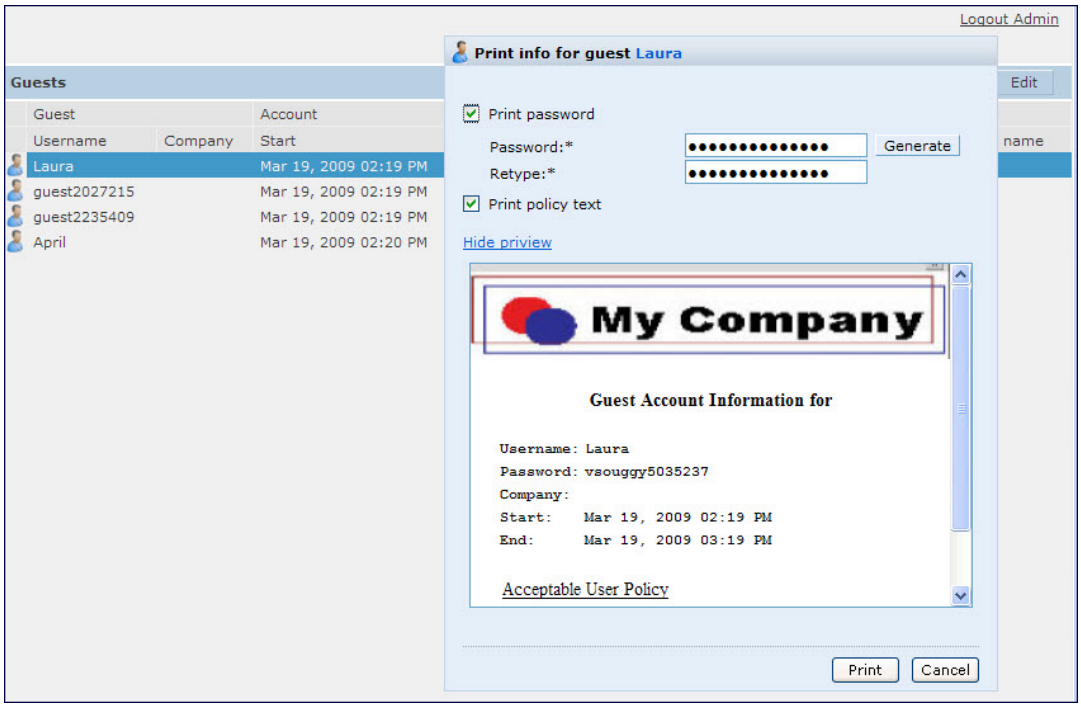


Printing Guest Account Information

To print guest account information:

1. Highlight the guest account you want to print and click **Print**. The **Print info for guest** window displays.
2. Click **Print password** if you want to print the guest password on the badge. Then enter or generate a new password for the guest. This modifies the existing guest password. (See Figure 177.)
3. Optionally, click **Print policy text** if you want your company policy text to appear on the print out.
4. Click **Show preview** to view the information before it is printed.
5. Click **Print** to print the guest account information.

Figure 177 *Printing Guest Account Information*



Optional Configurations

This section describes guest provisioning options that the administrator can configure.

NOTE

These options are not configurable by the guest provisioning user.

Restricting one Captive Portal Session for each Guest

You can restrict one captive portal session for each guest. When a new captive portal request is received and passes authentication, all users are checked and compared with user names. If a user with the same name already exists and this option is enabled, the second login is denied.

NOTE

If a guest logs in from one system (and does not log out) and tries to log in again from another system, that guest has to wait for the initial session to expire.

1. Navigate to the **Configuration > Advanced Services> All s** page.
2. Select **Wireless Lan**.
3. Under Wireless Lan, select and open **Captive Portal Authentication**.
4. Add a new or select and existing
5. Select the **Allow only one active user session** check box.
6. Click **Apply**.

Using the CLI to restrict one Captive Portal session for each guest

```
(host)(config)# aaa authentication captive-portal <> single-session
```

Setting the Maximum Time for Guest Accounts

You can set the maximum expiration time (in minutes) for guest accounts. If the guest-provisioning user attempt to add a guest account that expires beyond this time period, an error message is displayed and the

guest account is created with the maximum time you configured.



If you set the maximum expiration time, it applies to all users in the internal database whether they are guests or not.

Using the WebUI to set the maximum time for guest accounts

1. Navigate to the **Configuration > Security > Authentication** page.
2. Select **Internal DB**.
3. Under Internal DB Maintenance, enter a value in Maximum Expiration.
4. Click **Apply**.

Using the CLI to set the maximum time for guest accounts

```
(host)# local-userdb maximum-expiration <minutes>
```

Managing Files on the Controller

You can transfer the following types of files between the controller and an external server or host:

- ArubaOS image file
- A specified file in the controller's flash file system, or a compressed archive file that contains the entire content of the flash file system.



You back up the entire content of the flash file system to a compressed archive file, which you can then copy from the flash system to another destination.

- Configuration file, either the active running configuration or a startup configuration.
- Log files.

You can use the following protocols to copy files to or from a controller:

- File Transfer Protocol (FTP): Standard TCP/IP protocol for exchanging files between computers.
- Trivial File Transfer Protocol (TFTP): Software protocol that does not require user authentication and is simpler to implement and use than FTP.
- Secure Copy (SCP): Protocol for secure transfer of files between computers that relies on the underlying Secure Shell (SSH) protocol to provide authentication and security.



You can use SCP only for transferring image files to or from the controller, or transferring files between the flash file system on the controller and a remote host. The SCP server or remote host must support SSH version 2 protocol.

The following table lists the parameters that you configure to copy files to or from a controller.

Table 172: *File Transfer Configuration Parameters*

Server Type	Configuration
Trivial File Transfer Protocol (TFTP)	• tftphost - tftp host IPv4 / IPv6 address • filename - absolute path of filename • flash: - copy to the flash file system • destination: - destination file name • system: - system partition • partition - partition 0 / partition 1
File Transfer Protocol (FTP)	• ftphost - ftp server host name or IPv4/IPv6 address • username - user name to log into server • filename - absolute path of filename • system: - system partition • partition - partition 0 / partition 1
Secure Copy (SCP) You must use the CLI to transfer files with SCP.	• scphost - scp host of IPv4 / IPv6 address • username - user name to secur to log into the server • filename - absolute path of filename (otherwise, SCP searches for the file relative to the user's home directory) • flash: - copy to the flash file system • destfilename: - destination file name • system: - system partition • partition - partition 0 / partition 1

For example, you can copy an ArubaOS image file from an SCP server to a system partition on a controller or copy the startup configuration on a controller to a file on a TFTP server, You can also store the contents of a controller's flash file system to an archive file which you can then copy to an FTP server. You can use SCP to securely download system image files from a remote host to the controller or securely transfer a configuration file from flash to a remote host.

Transferring ArubaOS Image Files

You can download an ArubaOS image file onto a controller from a TFTP, FTP, or SCP server. In addition, the WebUI allows you to upload an ArubaOS image file from the local PC on which you are running the browser.

When you transfer an ArubaOS image file to a controller, you must specify the system partition to which the file is copied. The WebUI shows the current content of the system partitions on the controller. You have the option of rebooting the controller with the transferred image file.

In the WebUI

1. Navigate to the **Maintenance > Controller > Image Management** page.
2. Select TFTP, FTP, SCP, or Upload Local File.
3. Enter or select the appropriate values for the file transfer method.
4. Select the system partition to which the image file is copied.
5. Specify whether the controller is to be rebooted after the image file is transferred, and whether the current configuration is saved before the controller is rebooted.
6. Click **Upgrade**.

In the CLI

```
copy tftp: <tftphost> <filename> system: partition [0|1]}
copy ftp: <ftphost> <user> <filename> system: partition {0|1}
copy scp: <scphost> <username> <filename> system: partition [0|1]
```

Backing Up and Restoring the Flash File System

You can store the entire content of the flash file system on a controller to a compressed archive file. You can then copy the archive file to an external server for backup purposes. If necessary, you can restore the backup file from the server to the flash file system.

Backup the Flash File System in the WebUI

1. Navigate to the **Maintenance > File > Backup Flash** page.
2. Click **Create Backup** to back up the contents of the flash system to the flashbackup.tar.gz file.
3. Click **Copy Backup** to enter the Copy Files page where you can select the destination server for the file.
4. Click **Apply**.

Backup the Flash File System in the CLI

```
backup flash
copy flash: flashbackup.tar.gz tftp: <tftphost> <destfilename>
copy flash: flashbackup.tar.gz scp: <scphost> <username> <destfilename>
```

Restore the Flash File System in the WebUI

1. Navigate to the **Maintenance > File > Copy Files** page.
   a. For Source Selection, specify the server to which the flashbackup.tar.gz file was previously copied.
   b. For Destination Selection, select Flash File System.
   c. Click **Apply**.
2. Navigate to the **Maintenance > File > Restore Flash** page.
3. Click **Restore** to restore the flashbackup.tar.gz file to the flash file system.
4. Navigate to the **Maintenance > Switch > Reboot Switch** page.
5. Click **Continue** to reboot the controller.

Restore the Flash File System in the CLI

```
copy tftp: <tftphost> <srcfilename> flash: flashbackup.tar.gz
copy scp: <scphost> <username> <srcfilename> flash: flashbackup.tar.gz
restore flash
```

Copying Log Files

You can store log files into a compressed archive file which you can then copy to an external TFTP or SCP server. The WebUI allows you to copy the log files to a WinZip folder which you can display or save on your local PC.

In the WebUI

1. Navigate to the **Maintenance > File > Copy Logs** page.
2. For Destination, specify the TFTP or FTP server to which log files are copied.
3. Select Download Logs to download the log files into a WinZip file on your local PC,
4. Click **Apply**.

In the CLI

```
tar logs
copy flash: logs.tar tftp: <tftphost> <destfilename>
copy flash: logs.tar scp: <scphost> <username> <destfilename>
```

Copying Other Files

The flash file system contains the following configuration files:

- startup-config: Contains the configuration options that are used the next time the controller is rebooted. It contains all options saved by clicking the **Save Configuration** button in the WebUI or by entering the **`write memory`** CLI command. You can copy this file to a different file in the flash file system or to a TFTP server.
- running-config: Contains the current configuration, including changes which have yet to be saved. You can copy this file to a different file in the flash file system, to the startup-config file, or to a TFTP or FTP server.

You can copy a file in the flash file system or a configuration file between the controller and an external server.

In the WebUI

1. Navigate to the **Maintenance > File > Copy Files** page.
2. Select the source where the file or image exists.
3. Select the destination to where the file or image is to be copied.
4. Click **Apply**.

In the CLI

```
copy startup-config flash: <filename>
copy startup-config tftp: <tftphost> <filename>

copy running-config flash: <filename>
copy running-config ftp: <ftphost> <user> <filename> [<remote-dir>]
copy running-config startup-config
copy running-config tftp: <tftphost> <filename>
```

Setting the System Clock

You can set the clock on a controller manually or by configuring the controller to use a Network Time Protocol (NTP) server to synchronize its system clock with a central time source.

Manually Setting the Clock

You can use either the WebUI or CLI to manually set the time on the controller's clock.

In the WebUI

1. Navigate to the **Configuration > Management > Clock** page.
2. Under **Controller Date/Time**, set the date and time for the clock.
3. Under **Time Zone**, enter the name of the time zone and the offset from Greenwich Mean Time (GMT).
4. To adjust the clock for daylight savings time, click **Enabled** under Summer Time. Additional fields appear that allow you to set the offset from UTC, and the start and end recurrences.
5. Click **Apply**.

In the CLI

To set the date and time, enter the following command in privileged mode:

```
clock set <year> <month> <date> <hour> <minutes> <seconds>
```

To set the time zone and daylight savings time adjustment, enter the following commands in configure mode:

```
clock timezone <WORD> <-23 - 23>

clock summer-time <zone> [recurring]
```

```
<1-4> <start day> <start month> <hh:mm>
first <start day> <start month> <hh:mm>
last <start day> <start month> <hh:mm>
<1-4> <end day> <end month> <hh:mm>
first <end day> <end month> <hh:mm>
last <end day> <end month> <hh:mm>
[<-23 - 23>]
```

Clock Synchronization

You can use NTP to synchronize the controller to a central time source. Configure the controller to set its system clock using NTP by configuring one or more NTP servers.

For each NTP server, you can optionally specify the NTP iburst mode for faster clock synchronization. The iburst mode sends up ten queries within the first minute to the NTP server. (When iburst mode is not enabled, only one query is sent within the first minute to the NTP server.) After the first minute, the iburst mode typically synchronizes the clock so that queries need to be sent at intervals of 64 seconds or more.



The iburst mode is a configurable option and not the default behavior for the controller, as this option is considered "aggressive" by some public NTP servers. If an NTP server is unresponsive, the iburst mode continues to send frequent queries until the server responds and time synchronization starts.

In the WebUI

1. Navigate to the **Configuration > Management > Clock** page.
2. Under NTP Servers, click **Add**.
3. Enter the IP address of the NTP server.
4. Select (check) the iburst mode, if desired.
5. Click **Add**.

In the CLI

```
ntp server ipaddr [iburst]
```

Configuring NTP Authentication

The Network Time Protocol adds security to an NTP client by authenticating the server before synchronizing the local clock. NTP authentication works by using a symmetric key which is configured by the user. The secret key is shared by both the controller and an external NTP server. This helps identify secure servers from fraudulent servers.

In the WebUI

1. Navigate to the **Configuration > Management > Clock** page.
2. Under **NTP Authentication**, make sure **Enable** is selected. Enable is the default option.
3. Under **NTP Servers**, enter the NTP server IP address (IPv4/IPv6) in the **Server IP** field.
4. Under **NTP Identification Keys**, enter an identification key (a number between 1 and 65535) in the **Identification Key** field. Then add a secret string in the **MD5 Secret** field. The MD5 ID key must be an ASCII string up to 31 characters.
5. Click **Add**.
6. The identification key along with its corresponding MD5 secret string is displayed in the **Identification Keys** section.
7. Under **NTP Trusted Keys**, enter a string in the **Trusted Key** field. This is a subset of keys which are trusted. The trusted key value must be numeric values between 1 to 65535.

8. Click **Apply**.

In the CLI

This example enables NTP authentication, add authentication secret keys into the database, and specifies a subset of keys which are trusted. It also enables the iburst option.

```
(host) (config) #ntp authenticate
(host) (config) #ntp authentication-key <key-id> md5 <key-secret>
(host) (config) #ntp trusted-key <key-id>
(host) (config) #ntp server <ipaddr> <iburst> <key>
(host) (config) #ntp server <server IP> <iburst key> <key>
```

Timestamps in CLI Output

The timestamp feature can include a timestamp in the output of each show command issued in the command-line interface, indicating the date and time the command was issued. Note that the output of **show clock** and **show log** do not include timestamps, even when this feature is enabled.

To enable this feature, access the command-line interface in config mode and issue the command **clock append**.

```
(host) (config) #clock append
```

ClearPass Profiling with IF-MAP

This feature is used in conjunction with ClearPass Policy Manager. It sends HTTP User Agent Strings and mDNS broadcast information to ClearPass so that it can make more accurate decisions about what types of devices are connecting to the network.

In the WebUI

To enable and configure this feature:

1. Navigate to **Configuration >All Profiles>Other Profiles**.
2. Click the **CPPM IF-MAP** profile.
3. Configure this profile according to the following parameters:

Table 173: *CPPM IF-Map Configuration Parameters*

Parameter	Description
CPPM IF-Map Interface	Enables the feature
Host IP address	IP address or hostname of the CPPM IF-MAP server
Username	Username for the user who performs actions on the CPPM IF-MAP server. Range must be between 1-255 bytes in length.
Password	Password of the user who performs actions on the CPPM IF-MAP server.Range between 6-100 bytes in length.

In the CLI

To configure this feature using the CLI:

```
(host) (config) #ifmap
(host) (config) #ifmap cppm
```

```
(host) (CPPM IF-MAP Profile) #server host <host>
(host) (CPPM IF-MAP Profile) #port <port>
(host) (CPPM IF-MAP Profile) #passwd <psswd>
(host) (CPPM IF-MAP Profile) #enable
```

This show command show if the CCPM interface is enable and the CPPM server IP address, username and password.

```
(host) (CPPM IF-MAP Profile) #show ifmap cppm
CPPM IF-MAP Profile
-------------------
Parameter              Value
---------              -----
CPPM IF-MAP Interface  Enabled
CPPM IF-MAP Server     10.4.191.32:443 admin/********
```

This show command shows if state of all enabled CPPM servers.

```
(host) (CPPM IF-MAP Profile) #show ifmap state cppm
CPPM IF-MAP Connection State [Interface: Enabled]
-------------------------------------------------
Server           State
------           -----
10.4.191.32:443  UP
```

Whitelist Synchronization

ArubaOS allows controllers to synchronize their remote AP whitelists with the Dell Activate cloud-based services. When you configure Activate whitelist synchronization, the controller will securely contact the Activate server and download the contents of the whitelist on the Activate server to the whitelist on the controller. The controller and the Activate server must have layer-3 connectivity to communicate.

By default, this feature will both add new remote AP entries to the controller whitelist and delete any obsolete entries on the controller whitelist that were not on the Activate server whitelist. Select the add-only option to allow this feature to add or modify entries, but not delete any existing entries.

In the WebUI

To enable this feature using the WebUI,

1. Navigate to **Configuration>Network>Controller>Sync Whitelist Service**.
2. Select **Enable sync service**.
3. In the **Activate user** field, enter the user name for your Activate account.
4. In the **Activate password** field, enter the password for your Activate account.
5. (Optional) Click the **Frequency** drop-down list and configure how frequently the controller should synchronize its remote AP whitelist with the whitelist on the Activate server.
6. Click **Apply** to save your settings.

In the CLI

The following example enables the Activate whitelist service on the controller. The **add-only** parameter allows only the addition of entries to the Activate remote AP whitelist database. This parameter is enabled by default. If this setting is disabled, the activate-whitelist-download command can both add and remove entries from the Activate database.

```
(host)(config)# activate-service-whitelist
(host)(activate-service-whitelist) #username user2 password pA$$w0rd whitelist-enable
(host)(activate-service-whitelist)add-only
```

The following command is available in enable mode, and prompts the controller to synchronize its remote AP whitelist with the associated whitelist on the Activate server:

```
(host)# activate whitelist download
```

Downloadable Regulatory Table

The downloadable regulatory table feature allows for the update of country domain options without upgrading the ArubaOS software version. A separate file, called the Regulatory-Cert, containing AP regulatory information will be released periodically on **download.dell-pcw.com**. The Regulatory-Cert file can then be uploaded to a controller and pushed to deployed APs.

The Regulatory-Cert includes the following information for each AP:

- All countries supported in the current release of ArubaOS (not just United States or Rest of World or any subset of countries)
- Allowed channels for each country
- Max EIRP for each channel and each country in the allowed list. The max values are specified for each PHY-type at which the AP is allowed to transmit on. The classified PHY-types are
  - 802.11 OFDM rates (802.11a/g mode)
  - 802.11b rates (CCK rates)
  - 802.11n HT20 and 802.11ac VHT20 rates (MCS0-7)
  - 802.11n HT20 and 802.11ac VHT40(MCS0-7)
  - 802.11ac VHT80 rates
- DFS functionality for each channel and each country in the allowed list

Important Points to Remember

- When a Regulatory-Cert is activated, the new file is checked against the default file built into ArubaOS. If the file is of a newer version, the activation is allowed. If the file is of a lower version, then the activation is not completed. The controller's CLI displays the following message upon failure:

  ```
  (host0) #ap regulatory activate reg-data-1.0_00002.txt
  Failed to activate regulatory file reg-data-1.0_00002.txt. File Version should be greater than 1.0_43859
  ```

- APs do not rebootstrap or reboot on activation.
- If there is change in channel list or power level, APs will change the channel/power level. Impact is same as that of ARM channel/power change in that case.
- Clients are not disconnected upon regulatory file activation. Max latency impact during activation (with no channel changes) is less than 1s (applies to power change too).
- With channel change, the impact is similar to ARM channel change (depends on client behavior and if CSA is enabled or not).
- If support for the AP (Country) is added, the AP will move from AM to AP mode (if the AP is configured in AP mode of operation).

Copying the Regulatory-Cert

You can use the following protocols to copy the regulatory file to a controller:

- FTP
- TFTP

- SCP

Additionally, regulatory files saved to a USB drive can be uploaded to a controller equipped with a USB port.

You can copy the Regulatory-Cert to the controller using the WebUI or CLI.

In the WebUI

1. Navigate to the **Maintenance > File > Copy Files** page.
2. Select the source (TFTP, FTP, SCP, or USB) where the file exists.
3. The controller WebUI will automatically select **Flash File System** under the **Destination Selection** menu.
4. Click **Apply**.

In the CLI

Use one of the following **copy** commands to download the regulatory file to the controller:

```
copy
   ftp: <ftphost> <user> <filename>
   scp: <scphost> <username> <filename> flash: <destfilename>
   tftp: <tftphost> <filename> flash: <destfilename>
   usb: partition <partition-number> <filename> flash:  <destfilename>
```

To view the current regulatory and the content of the file, use the following commands:

```
show ap regulatory
show ap allowed-channels country-code <country-code> ap-type <ap-type
show ap allowed-max-eirp ap-name <ap-name> country-code <country-code>
show ap debug received-reg-table ap-name <ap-name>
```

Activating the Regulatory-Cert

Once the Regulatory-Cert has been added to the controller, the new regulatory information must be activated and pushed to the APs.

In the WebUI

To activate a specific regulatory file using the WebUI:

1. Navigate to **Maintenance > File > Regulatory Files**.
2. Select a regulatory file from the **File List**.
3. Click **Activate**.

In the CLI

To activate a specific regulatory file loaded on the controller, use the following command:

```
ap regulatory activate <filename>
```

To return to the factory default regulatory-cert, use the following command:

```
ap regulatory reset
```

In a master-local-standby deployment, the file syncing profile can be enabled to ensure that the regulatory-cert that is stored on the master is shared with its subordinate controllers. File syncing is enabled by default, with a default sync time of 30 minutes. The sync time can be set between 30 to 180 minutes, To configure the file syncing profile, use the following commands

```
(host) (config) #file syncing profile
(host) (File syncing profile) #file-syncing-enable
(host) (File syncing profile) #sync-time 30
```

Related Show Commands

To view the version of Regulatory Cert currently active on all controllers, execute the following command:

```
(host) #show switches regulatory
```

To view the file synching profile settings, execute the following command:

```
(host) #show file syncing profile
```

ArubaOS incorporates Passpoint technology from the Wi-Fi Alliance Hotspot 2.0 Specification to simplify and automate access to public Wi-Fi networks. Follow the procedures in this chapter to help mobile devices identify which access points in your hotspot network are suitable for their needs, and authenticate to a remote service provider using suitable credentials.

Hotspot 2.0 Pre-Deployment Information

Hotspot 2.0 is a Wi-Fi Alliance Passpoint specification based upon the 802.11u protocol that provides wireless clients with a streamlined mechanism to discover and authenticate to suitable networks, and allows mobile users the ability to roam between partner networks without additional authentication. For an overview Hotspot 2.0 enhanced network discovery and selection technology, and a description of each of the hotspot profile types, see Hotspot 2.0 Overview on page 831

Hotspot Profile Configuration Tasks

The following sections describe the procedure to configure the profiles for the hotspot feature.

- Configuring Hotspot 2.0 Profiles on page 834
- Configuring Hotspot Advertisement Profiles on page 839
- Configuring ANQP Venue Name Profiles on page 840
- Configuring ANQP Network Authentication Profiles on page 842
- Configuring ANQP Domain Name Profiles on page 843
- Configuring ANQP IP Address Availability Profiles on page 844
- Configuring ANQP NAI Realm Profiles on page 845
- Configuring ANQP Roaming Consortium Profiles on page 848
- Configuring ANQP 3GPP Cellular Network Profiles on page 849
- Configuring H2QP Connection Capability Profiles on page 850
- Configuring H2QP Operator Friendly Name Profiles on page 852
- Configuring H2QP Operating Class Indication Profiles on page 853
- Configuring H2QP WAN Metrics Profiles on page 853

Hotspot 2.0 Overview

ArubaOS supports Hotspot 2.0 with enhanced network discovery and selection. Clients can receive general information about the network identity, venue and type via management frames from the Dell AP. Clients can also query APs for information about the network's available IP address type (IPv4 or IPv6), roaming partners, and supported authentication methods, and receive that information in Information Elements from the AP.

Generic Advertisement Service (GAS) Queries

An Organization Identifier (OI) is a unique identifier assigned to a service provider when it registers with the IEEE registration authority. A Dell AP can include its service provider OI in beacons and probe responses to clients. If a client recognizes an AP's OI, it will attempt to associate to that AP using the security credentials corresponding to that service provider.

If the client does not recognize the AP's OI, that client can send a Generic Advertisement Service (GAS) query to the AP to request more information more about the network before associating.

ANQP Information Elements

ANQP Information Elements (IEs) are additional data that can be sent from the AP to the client to identify the AP's network and service provider. If a client requests this information via a GAS query, the hotspot AP then sends the ANQP Capability list in the GAS Initial Response frame indicating support for the following IEs. If the client responds with a request for a specific IE, the AP will send a GAS response frame with the configured ANQP IE information.

- Venue Name: the Venue Name IE defines the venue group and venue type.
- Domain Name: this IE specifies the AP's domain name.
- Network Authentication Type: if the network has Additional Steps required for Access (ASRA), this profile defines the authentication type being used by the hotspot network.
- Roaming Consortium List: roaming Consortium Information Elements (IEs) contain information identifying the network and service provider, whose security credentials can then be used to authenticate with the AP transmitting this element.
- IP address Availability: this IE provides clients with information about the availability of IP address versions and types which could be allocated to those clients after they associate to the hotspot AP.
- NAI Realm: an AP's NAI Realm profile identifies and describes a NAI realm accessible using the AP, and the method that this NAI realm uses for authentication.
- 3GPP Cellular Network Data: defines information for a 3rd Generation Partnership Project (3GPP) Cellular Network for hotspots that have roaming relationships with cellular operators.
- Connection Capability: define hotspot protocol and port capabilities to be sent in an ANQP IE.
- Operating Class: use this profile to define the channels on which the hotspot is capable of operating.
- Operator Friendly Name: a free-form text field that can identify the operator and also something about the location.
- WAN Metrics: provides hotspot clients information about access network characteristics such as link status and the capacity and speed of the WAN link to the Internet.

Hotspot Profile Types

ArubaOS supports several different ANQP and H2QP profile types for defining Hotspot data. The following table describes the profiles in the Hotspot profile set.

Table 174: *ANQP and H2QP Profiles referenced by an Advertisement Profile*

Profile	Description
Hotspot Advertisement profile	An advertisement profile defines a collection of ANQP and H2QP profiles. Each hotspot 2.0 profile is associated with one advertisement profile, which in turn references one of each type of ANQP and H2QP profile. For more information on configuring this profile, refer to Configuring Hotspot Advertisement Profiles on page 839
ANQP 3GPP Cellular Network profile	Use this profile to define priority information for a 3rd Generation Partnership Project (3GPP) Cellular Network used by hotspots that have roaming relationships with cellular operators. For more information on configuring this profile, refer to Configuring ANQP 3GPP Cellular Network Profiles on page 849
ANQP Domain Name profile	Use this profile to specify the hotspot operator domain name. For more information on configuring this profile, refer to Configuring Hotspot Advertisement Profiles on page 839
ANQP IP Address Availability profile	Use this profile to specify the types of IPv4 and IPv6 IP addresses available in the hotspot network. For more information on configuring this profile, refer to Configuring ANQP IP Address Availability Profiles on page 844
ANQP NAI Realm profile	An AP's NAI Realm profile identifies and describes a Network Access Identifier (NAI) realm accessible using the AP, and the method that this NAI realm uses for authentication. For more information on configuring this profile, refer to Configuring ANQP NAI Realm Profiles on page 845
ANQP Network Authentication profile	Use the ANQP Network Authentication profile to define the authentication type used by the hotspot network. For more information on configuring this profile, refer to Configuring ANQP Network Authentication Profiles on page 842.
ANQP Roaming Consortium profile	Name of the ANQP Roaming Consortium profile to be associated with this WLAN advertisement profile. For more information on configuring this profile, refer to Configuring ANQP Roaming Consortium Profiles on page 848
ANQP Venue Name profile	Use this profile to specify the venue group and venue type information be sent in an Access network Query Protocol (ANQP) information element in a Generic Advertisement Service (GAS) query response. For more information on configuring this profile, refer to Configuring ANQP Venue Name Profiles on page 840.
H2QP Connection Capability profile	Use this profile to specify hotspot protocol and port capabilities. For more information on configuring this profile, refer to Configuring H2QP Connection Capability Profiles on page 850

Profile	Description
H2QP Operating Class Indication profile	Use this profile to specify the channels on which the hotspot is capable of operating For more information on configuring this profile, refer to Configuring H2QP Operating Class Indication Profiles on page 853
H2QP Operator Friendly Name profile	Use this profile to define the operator-friendly name sent by devices using this profile. For more information on configuring this profile, refer to Configuring H2QP Operator Friendly Name Profiles on page 852
H2QP WAN Metrics profile	Use this profile to specify the WAN status and link metrics for your hotspot. For more information on configuring this profile, refer to Configuring H2QP WAN Metrics Profiles on page 853

Configuring Hotspot 2.0 Profiles

Use this profile to enable the hotspot 2.0 feature, and define venue and OI settings for roaming partners. Each hotspot 2.0 profile also references an advertisement profile, which defines a set of ANQP or H2QP profiles which define other values for the hotspot feature. By default, hotspot 2.0 profiles references the **default** advertisement profile. For information on associating a different advertisement profile with a hotspot 2.0 profile, see Associating the Advertisement Profile to a Hotspot 2.0 Profile on page 840.

In the WebUI

To configure a hotspot 2.0 profile from the controller WebUI:

1. Navigate to **Configuration>Advanced Services>All Profiles**.
2. In the profiles list, expand the **Wireless LAN** section.
3. Select **Hotspot 2.0**.
4. Select an existing profile from the list of profiles on the **profile details** pane or create a new profile by entering a profile name into the entry blank, then clicking **Add**.
5. Configure the parameters described in Table 175 as desired, then click **Apply.**

Table 175: *Hotspot 2.0 Profile Settings*

Parameter	Description
Advertise Hotspot 2.0 Capability	This checkbox enables or disables the hotspot. When this feature is enabled, the Information Elements (IEs) for this hotspot are included in beacons and probe responses from the AP. This setting is disabled by default.
Use GAS Comeback Request/Response	By default, ANQP Information is obtained from a GAS Request and Response. If you enable the Use GAS Comeback Request/Response option, advertisement information is obtained using a GAS Request and Response. as well as a Comeback-Request and Comeback-Response. This option is disabled by default.
Additional Steps required for Access Enabled	Select this checkbox if any additional steps are required for network access. If this parameter is enabled, the AP will send the following Information Elements (IEs) in response to the client's the ANQP query. • Venue Name

Parameter	Description
	• Domain Name List • Network Authentication Type • Roaming Consortium List • NAI Realm List **NOTE:** If this parameter is enabled, the advertisement profile for this hotspot must reference an enabled network authentication type profile.
Network Internet Access	If you select this checkbox, the AP sends an Information Element (IE) indicating that the network allows internet access. By default, a hotspot profile does not advertise network internet access.
Length of Query Response	Generic Advertisement Service (GAS) enables advertisement services that lets clients query multiple 802.11 networks at once, while also allowing the client to learn more about a network's 802.11 infrastructure before associating. If a client transmits a GAS Query using a GAS Initial Request frame, the responding AP will provide the query response (or information on how to receive the query response) in a GAS Initial Response frame. This parameter sets the maximum length of the GAS query response, in octets. The supported range is 1-255 octets.
Access Network Type	Specify the 802.11u network type. The default setting is public-chargeable. • **emergency-services**: emergency services only network • **personal-device**: personal device network • **private**: private network • **private-guest**: private network with guest access • **public-chargeable**: public chargeable network • **public-free**: free public network • **test**: test network • **wildcard**: wildcard network
Roaming Consortium Len Entry 1	Length of the OI. The value of the Roaming Consortium Len Entry 1 parameter is based upon the number of octets of the Roaming Consortium OI Entry 1 field. • **0**: Zero Octets in the OI (Null) • **3**: OI length is 24-bit (3 Octets) • **5**: OI length is 36-bit (5 Octets)
Roaming Consortium OI Entry 1	Roaming consortium OI assigned to one of the service provider's top three roaming partners.This additional OI will only be sent to a client if the Additional Roaming Consortium OI's parameter is set to1 or higher. **NOTE:** The service provider's own roaming consortium OI is configured using the ANQP Roaming Consortium profile.
Roaming Consortium Len Entry 2	Length of the OI. The value of the Roaming Consortium Len Entry 2 parameter is based upon the number of octets of the Roaming Consortium OI Entry 2 field. • **0**: Zero Octets in the OI (Null) • **3**: OI length is 24-bit (3 Octets) • **5**: OI length is 36-bit (5 Octets)

Parameter	Description
Roaming Consortium OI Entry 2	Roaming consortium OI assigned to one of the service provider's top three roaming partners. This additional OI will only be sent to a client if the Additional Roaming Consortium OI's parameter is set to 2 or higher. **NOTE:** The service provider's own roaming consortium OI is configured using the ANQP Roaming Consortium profile.
Roaming Consortium Len Entry 3	Length of the OI. The value of the Roaming Consortium Len Entry - 3 parameter is based upon the number of octets of the Roaming Consortium OI Entry 3 field. • **0**: Zero Octets in the OI (Null) • **3**: OI length is 24-bit (3 Octets) • **5**: OI length is 36-bit (5 Octets)
Roaming Consortium OI Entry 3	Roaming consortium OI assigned to one of the service provider's top three roaming partners. This additional OI will only be sent to a client if the Additional Roaming Consortium OI's parameter is set to 3. **NOTE:** The service provider's own roaming consortium OI is configured using the ANQP Roaming Consortium profile.
Additional Roaming Consortium OI's (displayed in Advertisement Profile)	Number of additional roaming consortium Organization Identifiers (OIs) advertised by the AP. This feature supports up to three additional OIs, which are defined using the Roaming Consortium Len Entry 1, Roaming Consortium Len Entry 2 and Roaming Consortium Len Entry 3 parameters
HESSID	This optional parameter devices an AP's homogenous ESS identifier (HESSSID), which is that device's MAC address in colon-separated hexadecimal format.
Venue Group Type	Specify one of the following venue groups to be advertised in the IEs from APs associated with this hotspot profile. The default setting is unspecified. • assembly • business • educational • factory-or-industrial • institutional • mercantile • outdoor • reserved • residential • storage • unspecified • Utility-Misc • Vehicular **NOTE:** This parameter only defines the venue group advertised in the IEs from hotspot APs.
Venue Type	Specify a venue type to be advertised in the IEs from APs associated with this hotspot profile. The complete list of supported venue types is described in Configuring Hotspot 2.0 Profiles on page 834. This parameter only defines the venue type advertised in the IEs from hotspot APs.

Parameter	Description
PAME BI	This option enables the Pre-Association Message Exchange BSSID Independent (PAME-BI) bit, which is used by an AP to indicate whether the the AP indicates that the Advertisement Server can return a query response that is independent of the BSSID used for the GAS Frame exchange.
Downstream Group Frames Forwarding Blocked	This option configures the Downstream Group Addressed Forwarding (DGAF) Disabled Mode. If this feature is enabled, it ensures that the AP does not forward downstream group-addressed frames. It is disabled by default, allowing the AP to forward downstram group-addressed frames.
Time Zone Format	The time zone in which the AP is operating, in the format `<std><offset>[dst[offset][,start[/time],end[/time]]` Where the <std> string specifies the abbreviation of the time zone, <dst> is the abbreviation of the timezone in daylight savings time, and the <offset> string specifies the time value you must add to the local time to arrive at UTC. **NOTE:** For complete details on configuring the timezone format, refer to section 8.3 of IEEE Std 1003.1, 2004 Edition.
Time Advertisement Capability	This parameter specifies the AP's source of external time, and the current condition of its timing estimator. • **no-std-ext-time-src**: The AP using this profile has no standardized external time source. • **timestamp-offset-utc**: The AP has a timestamp offset based on UTC. • **reserved**: This setting is reserved for future use, and should not be used.
Time Error Value	The standard deviation of error in the time value estimate, in milliseconds. The default value is 0 milliseconds, and the supported range is 0- 2,147,483,647 milliseconds.
P2P Device Management	Issue this command to advertise support for P2P device management. This setting is disabled by default.
P2P Cross Connect	Issue this command to advertise support for P2P Cross Connections. This setting is disabled by default.
Hotspot 2.0 Advertisement Protocol Type	Select one of the following advertisement protocol types to be used by the AP. • **anqp**: Access Network Query Protocol (ANQP) • **emergency**: Emergency Alert System (EAS) • **mih-cmd-event**: Media Independent Handover (MIH) Command and Event Services Capability Discovery • **mih-info**: Media Independent Handover (MIH) Information Service. This option allows handovers between differing kinds of wireless access protocols and technologies, allowing access points on different IP subnets to communicate with each other at the link level while maintaining session continuity. • **rsvd**: Reserved for future use.

Parameter	Description
GAS comeback delay in milliseconds	At the end of the GAS comeback delay interval, the client may attempt to retrieve the query response using a Comeback Request Action frame. The supported range is 100-2000 milliseconds, and the default value is 500 milliseconds.
RADIUS Chargeable User Identity (RFC4372)	Include this parameter to enable the Chargeable-User-Identity RADIUS attribute defined by RFC 4372. Home networks can use this attribute to identify a user for the roaming transactions that take place outside of that home network.
RADIUS Location Data (RFC5580)	Include this parameter to enable the Location Data RADIUS attribute defined by RFC 5580. Enabling this parameter allows the RADIUS server to use user location data.

In the CLI

To configure a hotspot 2.0 profile from the controller CLI, access the CLI in config mode and issue the following commands:

```
wlan hotspot h2-profile <profile-name>
   access-network-type emergency-services|personal-device|private|private-guest|public-
   chargeable|public-free|test|wildcard
   addtl-roam-cons-ois <addtl-roam-cons-ois>
   advertisement-profile <profile-name>
   advertisement-protocol anqp|eas|mih-cmd-event|mih-info|rsvd
   asra
   clone <profile-name>
   comeback-mode
   gas-comeback-delay
   grp-frame-block
   hessid <id>
   hotspot-enable
   internet
   no ..
   p2p-cross-connect
   p2p-dev-mgmt
   pame-bi
   query-response-length-limit <query-response-length-limit>
   radius_cui
   radius_loc_data
   roam-cons-len-1 0|3|5
   roam-cons-len-2 0|3|5
   roam-cons-len-3 0|3|5
   roam-cons-oi-1 <roam-cons-oi-1>
   roam-cons-oi-2 <roam-cons-oi-1>
   roam-cons-oi-3 <roam-cons-oi-1>
   time-advt-cap no-std-ext-timesrc|timestamp-offset-utc |reserved
   time-error <milliseconds>
   time-zone <time-zone>
   venue-group <venue-group>
   venue-type <venue-type>
```

Configuring Hotspot Advertisement Profiles

An advertisement profile defines a set of ANQP and H2QP profiles for the hotspot feature. Advertisement profiles can reference multiple instances of some ANQP and H2QP profile types, but only a single instance of other ANQP and H2QP profiles. The table below shows how the different ANQP and H2QP profile types can be associated to a single advertisement profile.

Table 176: *Hotspot Advertisement Profile Associations*

One Instance per Advertisement Profile	Multiple Instances per Advertisement Profile
• ANQP IP address availability profile • H2QP WAN metrics profile • H2QP connection capability profile	• ANQP venue name profile • ANQP network authentication profile • ANQP foaming consortium profile • ANQP NAI realm profile • ANQP 3GPP cellular network profile • H2QP operator friendly name profile • H2QP operating class indication profile • ANQP domain Name profile



For more information on each of these profile types, see Hotspot Profile Types on page 832

Configuring an Advertisement Profile

The steps below describe the procedure to associate an advertisement profile to a set of ANQP and H2QP profiles. Note that the procedure to associate an ANQP or H2QP profile to an advertisement profile varies, depending upon whether the advertisement profile can reference just one instance or many instances of that profile type

In the WebUI

To configure an advertisement profile from the controller WebUI:

1. Navigate to **Configuration>Advanced Services>All Profiles**.
2. In the **Profiles** list, click **Wireless LAN** expand the Wireless LAN profiles section.
3. Select **Advertisement.**
4. Select an existing advertisement profile from the list of profiles in the **Profiles** list. pane or create a new advertisement profile by entering a profile name into the entry blank on the **Profile Details** pane, then clicking **Add**. The ANQP and H2QP profiles associated with the selected advertisement profile appear below the advertisement profile in the **Profiles** list.
5. For an ANQP or H2QP profile type that can have only one instance associated with the advertisement profile:
   a. In the **Profiles** list, select the ANQP or H2QP profile type.
   b. Click the drop-down list in the **Profile Details** pane and select a profile name.
6. For an ANQP or H2QP profile type that can have multiple instances associated with the advertisement profile:
   a. In the **Profiles** list, select the ANQP or H2QP profile type.
   b. In the **Profile Details** pane, click the **Add a Profile** drop down list.
   c. Select the name of the profile to associate with the advertisement profile.
   d. click **Add**.

e. (Optional) To remove an existing reference to an ANQP or H2QP profile, select the profile name in the **Profile Details** pane, then click **Delete**.

7. Click **Apply** .

In the CLI

To configure a advertisement profile from the controller CLI, access the CLI in config mode and issue the following commands:

```
wlan hotspot advertisement profile <profile-name>
   anqp-3gpp-nwk-profile <profile-name>
   anqp-domain-name-profile <profile-name>
   anqp-ip-addr-avail-profile <profile-name>
   anqp-nai-realm-profile <profile-name>
   anqp-nwk-auth-profile <profile-name>
   anqp-roam-cons-profile <profile-name>
   anqp-venue-name-profile <profile-name>
   clone <profile-name>
   h2qp-conn-cap-profile <profile-name>
   h2qp-op-cl-profile <profile-name>
   h2qp-operator-friendly-profile <profile-name>
   h2qp-wan-metrics-profile <profile-name>
   no ...
```

Associating the Advertisement Profile to a Hotspot 2.0 Profile

The settings in the ANQP and H2QP profiles referenced by the Advertisement profile will not be sent to clients until you associate the advertisement profile with an active hotspot 2.0 profile. By default, all hotspot 2.0 profiles reference the **default** advertisement profile.

In the WebUI

To associate a different advertisement profile to a hotspot 2.0 profile:

1. Navigate to **Configuration>Advanced Services>All Profiles**.
2. In the **Profiles** list, click **Wireless LAN** expand the Wireless LAN profiles section.
3. Select **Hotspot 2.0**. The list of available hotspot 2.0 profiles appears in the **Profiles** list.
4. In the **Profiles** list, select a hotspot 2.0 profile.
5. Click the **Advertisement** link that appears below the selected hotspot 2.0 profile.
6. In the **Profile Details** list, click the **Advertisement Profile** drop-down list and select a different advertisement profile name.
7. Click **Apply**.

In the CLI

To associate a different advertisement profile to a hotspot 2.0 profile from the controller CLI, access the CLI in config mode and issue the following commands:

```
wlan hotspot hs2-profile <hotspot-profile-name>
   advertisement-profile <advertisement-profile-name>
```

Configuring ANQP Venue Name Profiles

Use this profile to define the venue group and venue type information be sent in an Access network Query Protocol (ANQP) information element in a Generic Advertisement Service (GAS) query response. If a client uses the Generic Advertisement Service (GAS) to post an ANQP query to an Access Point, the AP will return ANQP Information Elements with the values configured in this profile.

To send the values configured in this profile to clients, you must associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For details, see Configuring Hotspot Advertisement Profiles on page 839.

In the WebUI

To configure an ANQP venue name profile from the controller WebUI:

1. Navigate to **Configuration>Advanced Services>All Profiles**.
2. In the profiles list, expand the **Wireless LAN** section.
3. Select **ANQP Venue Name.**
4. Select an existing profile from the list of profiles on the **profile details** pane or create a new profile by entering a profile name into the entry blank, then clicking **Add**.
5. Configure the following parameters as desired, then click **Apply** to save your settings.

Table 177: *ANQP Venue Name Profile Parameters*

Parameter	Description
Venue Group	Specify one of the following venue groups to be advertised in the ANQP Information Elements (IEs) from APs associated with this profile. The default setting is unspecified. • assembly • business • educational • factory-or-industrial • institutional • mercantile • outdoor • reserved • residential • storage • unspecified • Utility-Misc • Vehicular
Venue Language Code	An ISO 639 language code that identifies the language used in the Venue Name field.
Venue Name	Venue name to be advertised in the ANQP IEs from APs associated with this profile. If the venue name includes spaces, the name must be enclosed in quotation marks, e.g. “Midtown Shopping Center”.
Venue Type	Specify a venue type to be advertised in the IEs from APs associated with this hotspot profile. The complete list of supported venue types is described the table below

Venue Types

The following list describes the different venue types that may be configured in a Hotspot 2.0 or ANQP Venue Name profile:

• assembly-amphitheater • assembly-amusement-park • assembly-arena • assembly-bar • assembly-coffee-shop • assembly-convention-center • assembly-emer-coord-center • assembly-library • assembly-museum • assembly-passenger-terminal • assembly-restaurant • assembly-stadium • assembly-theater • assembly-worship-place • assembly-zoo • business-attorney • business-bank • business-doctor	• business-fire-station • business-police-station • business-post-office • business-professional-office • business-research-and-development • educational-primary-school • educational-secondary-school • educational-university • industrial-factory • institutional-alcohol-or-drug-rehab • institutional-group-home • institutional-hospital • institutional-prison • institutional-terminal-care • mercantile-automotive-service-station • mercantile-gas-station • mercantile-grocery • mercantile-retail	• mercantile-shopping-mall • outdoor-bus-stop • outdoor-city-park • outdoor-kiosk • outdoor-muni-mesh-nwk • outdoor-rest-area • outdoor-traffic-control • residential-boarding-house • residential-dormitory • residential-hotel • residential-private-residence • unspecified • vehicular-airplane • vehicular-automobile • vehicular-bus • vehicular-ferry • vehicular-motor-bike • vehicular-ship • vehicular-train

In the CLI

To configure an ANQP venue name profile from the controller CLI, access the CLI in config mode and issue the following commands:

```
wlan hotspot anqp-venue-name-profile <profile-name>
   clone <profile-name>
   no ...
   venue-group outdoor|reserved|utility-misc|vehicular|assembly|business  educational|factory-
   or-industrial|institutional|mercantile|residential|   storage|unspecified
   venue-language <language>
   venue-name <venue-name>
   venue-type <venue-type>
```

Configuring ANQP Network Authentication Profiles

Use the ANQP Network Authentication profile to define the authentication type used by the hotspot network.

To send the values configured in this profile to clients, you must associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For details, see Configuring Hotspot Advertisement Profiles on page 839.

In the WebUI

To configure an ANQP network authentication profile from the controller WebUI:

1. Navigate to **Configuration>Advanced Services>All Profiles**.
2. In the profiles list, expand the **Wireless LAN** section.

3. Select **ANQP Network Authentication.**
4. Select an existing profile from the list of profiles on the **profile details** pane or create a new profile by entering a profile name into the entry blank, then clicking **Add**.
5. Configure the following parameters as desired, then click **Apply** to save your settings.

Table 178: *ANQP Network Authentication Profile Parameters*

Parameter	Description
Type of Network Authentication	Network Authentication Type being used by the hotspot network. • **acceptance**: Network requires the user to accept terms and conditions. This option requires you to specify a redirection URL string as an IP address, FQDN or URL. • **dns-redirection**: Additional information on the network is provided through DNS redirection. This option requires you to specify a redirection URL string as an IP address, FQDN or URL. • **http-https-redirection**: Additional information on the network is provided through HTTP/HTTPS redirection. • **online-enroll**: Network supports online enrollment.
Network Authentication URL	URL, IP address, or FQDN used by the hotspot network for the **acceptance** or **dns-redirection** network authentication types.

In the CLI

To configure an ANQP network authentication profile from the controller CLI, access the CLI in config mode and issue the following commands:

```
wlan hotspot anqp-nwk-auth-profile <profile-name>
   clone <profile-name>
   no ...
   nwk-auth-type acceptance|dns-redirection|http-https-redirection|online-enroll
   url <url>
```

Configuring ANQP Domain Name Profiles

This profile defines the hotspot operator domain name to be sent in an Access Network Query Protocol (ANQP) information element in a Generic Advertisement Service (GAS) query response.

To send the values configured in this profile to clients, you must associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For details, see Configuring Hotspot Advertisement Profiles on page 839.

In the WebUI

To configure an ANQP domain name profile from the controller WebUI:

1. Navigate to **Configuration>Advanced Services>All Profiles**.
2. In the profiles list, expand the **Wireless LAN** section.
3. Select **ANQP Domain Name.**
4. Select an existing profile from the list of profiles on the **profile details** pane or create a new profile by entering a profile name into the entry blank, then clicking **Add**.
5. In the **Domain Name** field, enter the domain name of the hotspot operator. This alphanumeric text string must be 32 characters or less.
6. Click **Apply**.

In the CLI

To configure an ANQP domain name profile from the controller CLI, access the CLI in config mode and issue the following commands:

```
wlan hotspot anqp-domain-name-profile <profile-name>
   clone <profile-name>
   domain-name <domain-name>
   no ...
```

Configuring ANQP IP Address Availability Profiles

Use this profile to specify the types of IPv4 and IPv6 IP addresses available in the hotspot network. This information is sent in an Access network Query Protocol (ANQP) information element in a Generic Advertisement Service (GAS) query response.

To send the values configured in this profile to clients, you must associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For details, see Configuring Hotspot Advertisement Profiles on page 839.

In the WebUI

To configure an ANQP IP address availability profile from the controller WebUI:

1. Navigate to **Configuration>Advanced Services>All Profiles**.
2. In the profiles list, expand the **Wireless LAN** section.
3. Select **ANQP IP Address Availability**
4. Select an existing profile from the list of profiles on the **profile details** pane or create a new profile by entering a profile name into the entry blank, then clicking **Add**.
5. Configure the following parameters as desired, then click **Apply** to save your settings.

Table 179: *ANQP IP Address Availability Profile Parameters*

Parameter	Description
IPv4 Address Availability Type	Indicate the availability of an IPv4 network by clicking the **IPv4 Address Availability Type** drop-down list and selecting one of the following options: • availability-unknown: Network availability cannot be determined. • not-available: Network is not available. • port-restricted: Some ports are restricted (e.g., the network blocks port 110 to restrict POP mail). • port-restricted-double-nated: Some ports are restricted and multiple routers perform network address translation. • port-restricted-single-nated: Some ports are restricted and a single router performs network address translation. • private-double-nated: Network is a private network with multiple routers doing network address translation. • private-single-nated: Network is a private network a single router doing network address translation. • public: Network is a public network.
IPv6 Address Availability Type	Indicate the availability of an IPv6 network by clicking the **IPv6 Address Availability Type** drop-down list and selecting one of the following options: • available: An IPv6 network is available. • availability-unknown: Network availability cannot be determined. • not-available: Network is not available.

In the CLI

To configure an ANQP IP address availability profile from the controller CLI, access the CLI in config mode and issue the following commands:

```
wlan hotspot anqp-ip-addr-avail-profile <profile-name>
   clone <profile-name>
   ipv4-addr-avail availability-unknown|not-available|port-restricted|port-restricted-double-
   nated|port-restricted-single-nated|private-double-nated|private-single-nated
   ipv6-addr-avail available|availability-unknown|not-available
   no ...
```

Configuring ANQP NAI Realm Profiles

An AP's NAI Realm profile identifies and describes a Network Access Identifier (NAI) realm accessible using the AP, and the method that this NAI realm uses for authentication. These settings configured in this profile determine the NAI realm elements that are included as part of a GAS Response frame.

To send the values configured in this profile to clients, you must associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For details, see Configuring Hotspot Advertisement Profiles on page 839.

In the WebUI

To configure an ANQP NAI Realm profile from the controller WebUI:

1. Navigate to **Configuration>Advanced Services>All Profiles**.
2. In the profiles list, expand the **Wireless LAN** section.
3. Select **ANQP NAI Realm**
4. Select an existing profile from the list of profiles on the **profile details** pane or create a new profile by entering a profile name into the entry blank, then clicking **Add**.

5. Configure the following parameters as desired, then click **Apply**.

Table 180: *ANQP NAI Realm Profile Parameters*

Parameter	Description
NAI Realm name	Name of the NAI realm. The realm name is often the domain name of the service provider.
NAI Realm Encoding	Issue this command if the NAI realm name is a UTF-8 formatted character string that is not formatted in accordance with IETF RFC 4282.
NAI Realm EAP Method	Select one of the options below to identify the EAP authentication method supported by the hotspot realm. • crypto-card: Crypto card authentication • eap-aka: EAP for Universal Mobile Telecommunications System (UMTS) Authentication and Key Agreement • eap-sim: EAP for GSM Subscriber Identity Modules • eap-tls: EAP-Transport Layer Security • eap-ttls: EAP-Tunneled Transport Layer Security • generic-token-card: EAP Generic Token Card (EAP-GTC) • identity: EAP Identity type • notification: The hotspot realm uses EAP Notification messages for authentication. • one-time-password: Authentication with a single-use password • peap: Protected Extensible Authentication Protocol • peap-mschapv2: Protected Extensible Authentication Protocol with Microsoft Challenge Handshake Authentication Protocol version 2
NAI Realm Authentication Param ID 1	Use the **NAI Realm Authentication Param ID 1** parameter to send the one of the following authentication methods for the primary NAI realm ID. • credential-type: The specified authentication ID uses credential authentication. • expanded-eap: The specified authentication ID uses the expanded EAP authentication method. • expanded-inner-eap: The specified authentication ID uses the expanded inner EAP authentication method. • inner-auth-eap: The specified authentication ID uses inner EAP authentication type. • non-eap-inner-auth: The specified authentication ID uses non-EAP inner authentication type.
NAI Realm Authentication Param Value 1	Use the **NAI Realm Authentication Param Value 1** parameter select an authentication value for the authentication method specified by the **NAI Realm Authentication Param ID 1** parameter. • cred-cert: Credential - Certificate • cred-hw-token: Credential - Hardware Token • cred-nfc: Credential - NFC • cred-none: Credential - None • cred-rsvd: Credential - Reserved • cred-sim: Credential - SIM • cred-soft-token: Credential - Soft Token • cred-user-pass: Credential - Username/password • cred-usim: Credential - USIM • cred-vendor-spec: Credential - Vendor-specific • eap-crypto-card: EAP Method - Crypto-card • eap-generic-token-card: EAP Method - Generic-Token-Card • eap-identity: EAP Method - Identity • eap-method-aka: EAP Method - AKA • eap-method-sim: EAP Method - SIM - GSM Subscriber Iden

Parameter	Description
	• eap-method-tls: EAP Method - TLS - Transport Layer Sec • eap-method-ttls: EAP Method - TTLS - Tunneled Transport Security • eap-notification: EAP Method - Notification • eap-one-time-password: EAP Method - One-Time-Password • eap-peap: EAP Method - PEAP • eap-peap-mschapv2: EAP Method - PEAP MSCHAP V2 • non-eap-chap: Non-EAP Method - CHAP • non-eap-mschap: Non-EAP Method - MSCHAP • non-eap-mschapv2: Non-EAP Method - MSCHAPv2 • non-eap-pap: Non-EAP Method - PAP • non-eap-rsvd: Non-EAP Method - Reserved for future use • reserved: Reserved for Future use
NAI Realm Authentication Param ID 2	Use the **NAI Realm Authentication ID Value 2** parameter to send the one of the following authentication methods for the secondary NAI realm ID. • credential-type: The specified authentication ID uses credential authentication. • expanded-eap: The specified authentication ID uses the expanded EAP authentication method. • expanded-inner-eap: The specified authentication ID uses the expanded inner EAP authentication method. • inner-auth-eap: The specified authentication ID uses inner EAP authentication type. • non-eap-inner-auth: The specified authentication ID uses non-EAP inner authentication type.
NAI Realm Authentication Param Value 2	Use the **NAI Realm Authentication Param Value 2** parameter select an authentication value for the authentication method specified by the **NAI Realm Authentication Param ID 2** parameter. • cred-cert: Credential - Certificate • cred-hw-token: Credential - Hardware Token • cred-nfc: Credential - NFC • cred-none: Credential - None • cred-rsvd: Credential - Reserved • cred-sim: Credential - SIM • cred-soft-token: Credential - Soft Token • cred-user-pass: Credential - Username/password • cred-usim: Credential - USIM • cred-vendor-spec: Credential - Vendor-specific • eap-crypto-card: EAP Method - Crypto-card • eap-generic-token-card: EAP Method - Generic-Token-Card • eap-identity: EAP Method - Identity • eap-method-aka: EAP Method - AKA • eap-method-sim: EAP Method - SIM - GSM Subscriber Iden • eap-method-tls: EAP Method - TLS - Transport Layer Sec • eap-method-ttls: EAP Method - TTLS - Tunneled Transport Security • eap-notification: EAP Method - Notification • eap-one-time-password: EAP Method - One-Time-Password • eap-peap: EAP Method - PEAP • eap-peap-mschapv2: EAP Method - PEAP MSCHAP V2 • non-eap-chap: Non-EAP Method - CHAP • non-eap-mschap: Non-EAP Method - MSCHAP • non-eap-mschapv2: Non-EAP Method - MSCHAPv2 • non-eap-pap: Non-EAP Method - PAP

Parameter	Description
	• non-eap-rsvd: Non-EAP Method - Reserved for future use • reserved: Reserved for Future use
NAI Home Realm	Mark the realm in this profile as the NAI Home Realm

In the CLI

To configure an ANQP NAI realm profile from the controller CLI, access the CLI in config mode and issue the following commands:

```
wlan hotspot anqp-nai-realm-profile <profile-name>
   clone <profile-name>
   nai-home-realm
   nai-realm-auth-id-1|nai-realm-auth-id-2 {credential-type|expanded-eap|expanded-inner-
   eap|inner-auth-eap|non-eap-inner-auth|tunneled-eap-credential-type}
   nai-realm-auth-value-1|nai-realm-auth-value-2 {cred-cert|cred-hw-token|cred-nfc|cred-
   none|cred-rsvd|cred-sim|cred-soft-token|cred-user-pass|cred-usim|cred-vendor-spec|eap-
   crypto-card|eap-generic-token-card|eap-identity|eap-method-aka|eap-method-sim|eap-method-
   tls|eap-method-ttls|eap-notification|eap-one-time-password|eap-peap|eap-peap-mschapv2|non-
   eap-chap|non-eap-mschap|non-eap-mschapv2|non-eap-pap|non-eap-rsvd|reserved}
   nai-realm-eap-method crypto-card|eap-aka|eap-sim|eap-tls|eap-ttls|generic-token-
     card|identity|notification|one-time-password|peap|peap-mschapv2
   nai-realm-encoding
   nai-realm-name <nai-realm-name>
   no ...
```

Configuring ANQP Roaming Consortium Profiles

Organization Identifiers (OIs) are assigned to service providers when they register with the IEEE registration authority. You can specify the OI for the hotspot's service provider in the ANQP Roaming Consortium profile using the ANQP Roaming Consortium Profile. The Hotspot 2.0 profile also allows you to define and send up to three additional roaming consortium OIs for the service provider's top three roaming partners.

To send the values configured in this profile to clients, you must associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For details, see Configuring Hotspot Advertisement Profiles on page 839.

In the WebUI

To configure an ANQP roaming consortium profile from the controller WebUI:

1. Navigate to **Configuration>Advanced Services>All Profiles**.
2. In the profiles list, expand the **Wireless LAN** section.
3. Select **ANQP Roaming Consortium.**
4. Select an existing profile from the list of profiles on the **profile details** pane or create a new profile by entering a profile name into the entry blank, then clicking **Add**.
5. Configure the following parameters as desired, then click **Apply** to save your settings.

Table 181: *ANQP Roaming Consortium Profile Parameters*

Parameter	Description
Roaming consortium OI Len	Length of the OI. The value of the **Roaming consortium OI Len** parameter must equal upon the number of octets of the **Roaming Consortium OI** field. • **0**: 0 Octets in the OI (Null) • **3**: OI length is 24-bit (3 Octets) • **5**: OI length is 36-bit (5 Octets)
Roaming Consortium OI	Send the specified roaming consortium OI in a GAS query response. The OI must be a hexadecimal number 3-5 octets in length.

In the CLI

To configure an ANQP roaming consortium profile from the controller CLI, access the CLI in config mode and issue the following commands:

```
wlan hotspot anqp-roam-cons-profile <profile-name>
   clone <profile-name>
   no ...
   roam-cons-oi <roam-cons-oi>
   roam-cons-oi-len <roam-cons-oi-len>
```

Configuring ANQP 3GPP Cellular Network Profiles

Use this profile to define priority information for a 3rd Generation Partnership Project (3GPP) Cellular Network used by hotspots that have roaming relationships with cellular operators.

To send the values configured in this profile to clients, you must associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For details, see Configuring Hotspot Advertisement Profiles on page 839.

In the WebUI

To configure an ANQP 3GPP cellular network profile from the controller WebUI:

1. Navigate to **Configuration>Advanced Services>All Profiles**.
2. In the profiles list, expand the **Wireless LAN** section.
3. Select **ANQP 3GPP Cellular Network.**
4. Select an existing profile from the list of profiles on the **profile details** pane or create a new profile by entering a profile name into the entry blank, then clicking **Add**.
5. Configure the following parameters as desired, then click **Apply** to save your settings.

Table 182: *ANQP 3GPP Cellular Network Profile Parameters*

Parameter	Description
3GPP PLMN1	The Public Land Mobile Networks (PLMN) value of the highest-priority network. The PLMN is comprised of a 12-bit Mobile Country Code (MCC) and the 12-bit Mobile Network Code (MNC).
3GPP PLMN2	The Public Land Mobile Networks (PLMN) value of the second-highest priority network. The PLMN is comprised of a 12-bit Mobile Country Code (MCC) and the 12-bit Mobile Network Code (MNC).
3GPP PLMN3	The Public Land Mobile Networks (PLMN) value of the third-highest priority network. The PLMN is comprised of a 12-bit Mobile Country Code (MCC) and the 12-bit Mobile Network Code (MNC).
3GPP PLMN4	The Public Land Mobile Networks (PLMN) value of the fourth-highest priority network. The PLMN is comprised of a 12-bit Mobile Country Code (MCC) and the 12-bit Mobile Network Code (MNC).
3GPP PLMN5	The Public Land Mobile Networks (PLMN) value of the fifth-highest priority network. The PLMN is comprised of a 12-bit Mobile Country Code (MCC) and the 12-bit Mobile Network Code (MNC).
3GPP PLMN6	The Public Land Mobile Networks (PLMN) value of the sixth-highest priority network. The PLMN is comprised of a 12-bit Mobile Country Code (MCC) and the 12-bit Mobile Network Code (MNC).

In the CLI

To configure an ANQP 3GPP network profile from the controller CLI, access the CLI in config mode and issue the following commands:

```
wlan hotspot anqp-3gpp-nwk-profile <profile-name>
   3gpp_plmn1    <3GPP PLMN1 data>
   3gpp_plmn2    <3GPP PLMN2 data>
   3gpp_plmn3    <3GPP PLMN3 data>
   3gpp_plmn4    <3GPP PLMN4 data>
   3gpp_plmn5    <3GPP PLMN5 data>
   3gpp_plmn6    <3GPP PLMN6 data>
   clone <profile-name>
   enable
   no ...
```

Configuring H2QP Connection Capability Profiles

Use this profile to specify hotspot protocol and port capabilities. This information is sent in a Access Network Query Protocol (ANQP) information element in a Generic Advertisement Service (GAS) query response.

To send the values configured in this profile to clients, you must associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For details, see Configuring Hotspot Advertisement Profiles on page 839.

In the WebUI

To configure a H2QP connection capability profile from the controller WebUI:

1. Navigate to **Configuration>Advanced Services>All Profiles**.
2. In the profiles list, expand the **Wireless LAN** section.

3. Select **H2QP Connection Capability.**
4. Select an existing profile from the list of profiles on the **profile details** pane or create a new profile by entering a profile name into the entry blank, then clicking **Add**.
5. Configure the following parameters as desired, then click **Apply** to save your settings.

Table 183: *ANQP Connection Capability Profile Parameters*

Parameter	Description
H2QP Connection Capability ICMP Port	Select this option to enable the ICMP port. (port 0)
H2QP Connection Capability FTP port (TCP Protocol)	Select this option to enable the FTP port. (port 20)
H2QP Connection Capability SSH port (TCP Protocol)	Select this option to enable the SSH port. (port 22)
H2QP Connection Capability HTTP port (TCP Protocol)	Select this option to enable the HTTP port. (port 80)
H2QP Connection Capability TLS VPN port (TCP Protocol)	H2QP Connection Capability TLS VPN port(TCP Protocol)
H2QP Connection Capability PPTP VPN port (TCP Protocol)	Select this option to enable the PPTP port used by IPSec VPNs. (port 1723)
H2QP Connection Capability VOIP port (TCP Protocol)	Select this option to enable the TCP VoIP port. (port 5060)
H2QP Connection Capability VOIP port (UDP Protocol)	Select this option to enable the UDP VoIP port. (port 5060)
H2QP Connection Capability IKEv2 port for IPSec VPN	Select this option to enable the IPsec VPN port. (ports 500, 4500 and 0)
H2QP Connection Capability May be used by IKEv2 port for IPSec VPN	Select this option to enable the IKEv2 port 4500.
H2QP Connection Capability IKEv2 port for IPSec VPN	Select this option to enable the IKEv2 port 500.
H2QP Connection Capability ESP port(Used by IPSec VPN)	Include this parameter to enable the Encapsulating Security Payload (ESP) port used by IPSec VPNs. (port 0)

In the CLI

To configure a H2QP connection capability profile from the controller CLI, access the CLI in config mode and issue the following commands:

```
wlan hotspot h2qp-conn-capability-profile <profile>
   clone <profile-name>
   esp
   icmp
```

```
no ...
tcp-ftp
tcp-http
tcp-pptp-vpn
tcp-ssh
tcp-tls-vpn
tcp-voip
udp-ike2-4500
udp-ike2-500
udp-ipsec-vpn
udp-voip
```

Configuring H2QP Operator Friendly Name Profiles

This profile defines an operator-friendly name sent by devices using this profile.

To send the values configured in this profile to clients, you must associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For details, see Configuring Hotspot Advertisement Profiles on page 839.

In the WebUI

To configure a H2QP operating class profile from the controller WebUI:

1. Navigate to **Configuration>Advanced Services>All Profiles**.
2. In the profiles list, expand the **Wireless LAN** section.
3. Select **H2QP Operator Friendly Name.**
4. Select an existing profile from the list of profiles on the **profile details** pane or create a new profile by entering a profile name into the entry blank, then clicking **Add**.
5. Configure the following parameters as desired, then click **Apply** to save your settings.

Table 184: *H2QP Operator Friendly Name Profile Parameters*

Parameter	Description
Operator Friendly Name Language Code	An ISO 639 language code that identifies the language used in the **Operator Friendly Name** field
Operator Friendly Name	An operator-friendly name sent by devices using this profile. The name can be up to 64 alphanumeric characters, and can include special characters and spaces. If the name includes quotation marks (“), include a backslash character (\) before each quotation mark. (e.g. \"example\")

In the CLI

To configure a H2QP operator friendly name profile from the controller CLI, access the CLI in config mode and issue the following commands:

```
wlan hotspot h2qp-operator-friendly-name-profile <profile>
   clone <profile-name>
   no ...
   op-fr-name <op-fr-name>
   op-lang-code <op-lang-code>
```

Configuring H2QP Operating Class Indication Profiles

The values configured in this H2QP Operating Class Indication profile list the channels on which the hotspot is capable of operating. It may be useful where, for instance, a mobile device discovers a hotspot in the 2.4 GHz band but finds it is dual-band and prefers the 5 GHz band.

To send the values configured in this profile to clients, you must associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For details, see Configuring Hotspot Advertisement Profiles on page 839.

In the WebUI

To configure a H2QP operating class indication profile from the controller WebUI:

1. Navigate to **Configuration>Advanced Services>All Profiles**.
2. In the profiles list, expand the **Wireless LAN** section.
3. Select **H2QP Operating Class Indication.**
4. Select an existing profile from the list of profiles on the **profile details** pane or create a new profile by entering a profile name into the entry blank, then clicking **Add**.
5. In the **H2QP Operating Class** field, enter a valid operating class value. (For a definition of these global operating classes refer to Table E-4 of IEEE Std 802.11-2012, Annex E.)
6. Click **Apply**.

In the CLI

To configure a H2QP operating class profile from the controller CLI, access the CLI in config mode and issue the following commands:

```
wlan hotspot h2qp-op-cl-profile <profile>
   clone <profile-name>
   op-cl <1-255>
```

Configuring H2QP WAN Metrics Profiles

Use this profile to specify the WAN status and link metrics for your hotspot.

To send the values configured in this profile to clients, you must associate this profile with an advertisement profile, then associate the advertisement profile with a hotspot 2.0 profile. For details, see Configuring Hotspot Advertisement Profiles on page 839.

In the WebUI

To configure an ANQP venue name profile from the controller WebUI:

1. Navigate to **Configuration>Advanced Services>All Profiles**.
2. In the profiles list, expand the **Wireless LAN** section.
3. Select **H2QP WAN Metrics.**
4. Select an existing profile from the list of profiles on the **profile details** pane or create a new profile by entering a profile name into the entry blank, then clicking **Add**.
5. Configure the following parameters as desired, then click **Apply** to save your settings.

Table 185: *H2QP WAN Metrics Profile Parameters*

Parameter	Description
H2QP WAN metrics link status	Define the status of the WAN Link by clicking the **H2QP WAN metrics link status** drop-down list, and selecting one of the following values. The default link status is **reserved**, which indicates that the link status is unknown or unspecified • **link down:** WAN link is down. • **link test**: WAN link is currently in a test state. • **link up**: WAN link is up. • **reserved**: This parameter is reserved by the Hotspot 2.0 specification, and cannot be configured. Default: reserved
H2QP WAN metrics symmetric WAN link	Select this checkbox to indicate that the WAN Link has same speed in both the uplink and downlink directions.
H2QP WAN metrics link at capacity	Select this checkbox to indicate that the WAN Link has reached its maximum capacity. If this parameter is enabled, no additional mobile devices will be permitted to associate to the hotspot AP.
WAN Metrics uplink speed	This parameter defines the current WAN uplink speed in Kbps. If no value is set, this parameter will show a default value of 0 to indicate that the uplink speed is unknown or unspecified. Range: 0 - 2147483647, Default: 0
WAN Metrics downlink speed	This parameter defines the current WAN backhaul downlink speed in Kbps. If no value is set, this parameter will show a default value of 0 to indicate that the downlink speed is unknown or unspecified. Range: 0 - 2147483647, Default: 0
WAN Metrics uplink load	This parameter defines the percentage of the WAN uplink that is currently utilized. If no value is set, this parameter will show a default value of 0 to indicate that the downlink speed is unknown or unspecified. Range: 0-100; Default: 0
WAN Metrics downlink load	This parameter defines the percentage of the WAN downlink that is currently in use. If no value is set, this parameter will show a default value of 0 to indicate that the downlink speed is unknown or unspecified. Range: 0-100; Default: 0
WAN Metrics load measurement duration	Duration over which the downlink load is measured, in tenths of a second. Range: 0-65535; Default: 0

In the CLI

To configure a H2QP WAN metrics profile from the controller CLI, access the CLI in config mode and issue the following commands:

```
wlan hotspot h2qp-wan-metrics-profile <profile-name>
   at-capacity
   clone <profile-name>
   downlink-load
   downlink-speed
   load-dur
   no ...
   symm-link
   uplink-load
   uplink-speed
```

```
wan-metrics-link-status link_down|link_test|link_up|reserved
```

This chapter explains how to expand your network by adding a local controller to a master controller configuration. Typically, this is the first expansion of a network with just one controller (which is a master controller). This chapter is a basic discussion of creating master-local controller configurations. More complicated multi-controller configurations are discussed in other chapters.

This chapter describes the following topics:

- Moving to a Multi-Controller Environment on page 859
- Configuring Local Controllers on page 857

Configuring Local Controllers

This section highlights the difference in configuration for both of these scenarios.

The steps involved in migrating from a single to a multi-controller environment are:

1. Configure the role of the local controller to local and specify the IP address of the master.
2. Configure the layer-2 / layer-3 settings on the local controller (VLANs, IP subnets, IP routes).
3. Configure as trusted ports the ports the master and local controller use to communicate with each other.
4. For those APs that need to boot off the local controller, configure the LMS IP address to point to the new local controller.
5. Reboot the APs that are already on the network, so that they now connect to the local controller.

These steps are explained below.

You configure the role of a controller by running the initial setup on an unconfigured controller, or by using the WebUI, Controller Wizard, or CLI on a previously-configured controller.

Using the Initial Setup

Initial setup can be done using the browser-based Setup Wizard or by accessing the initial setup dialog via a serial port connection. Both methods are described in the *Dell Networking W-Series ArubaOS Quick Start Guide* and are referred throughout this section as "initial setup".

The initial setup allows you to configure the IP address of the controller and its role, in addition to other operating parameters. You perform the initial setup the first time you connect to and log into the controller or whenever the controller is reset to its factory default configuration (after executing a **write erase**, **reload** sequence).

When prompted to enter the controller role in the initial setup, select or enter **local** to set the controller operational mode to be a local controller. You are then prompted for the master controller IP address. Enter the IP address of the master controller for the WLAN network. Enter the preshared key (PSK) that is used to authenticate communications between controllers.



You need to enter the same PSK on the master controller and on the local controllers that are managed by the master.

Using the Web UI

For a controller that is up and operating with layer-3 connectivity, configure the following to set the controller as local:

1. Navigate to the **Configuration > Network > Controller > System Settings** page.
2. Set the Controller Role to Local.
3. Enter the IP address of the master controller. If master redundancy is enabled on the master, this address should be the VRRP address for the VLAN instance corresponding to the IP address of the controller.
4. Enter the PSK that is used to authenticate communications between controllers.



You need to enter the same PSK on the master controller and on the local controllers that are managed by the master.

Using the CLI

For a controller that is up and operating with layer-3 connectivity, configure the following to set the controller as local:

```
masterip <ipaddr> ipsec <key>
```

Configuring Layer-2/Layer-3 Settings

Configure the VLANs, subnets, and IP address on the local controller for IP connectivity.

Verify connectivity to the master controller by pinging the master controller from the local controller.

Ensure that the master controller recognizes the new controller as its local controller. The local controller should be listed with type **local** in the **Monitoring > Network > All WLAN Controllers** page on the master. It takes about 4 – 5 minutes for the master and local controllers to synchronize configurations.

Configuring Trusted Ports

On the local controller, navigate to the **Configuration > Network > Ports** page and make sure that the port on the local controller connecting to the master is trusted. On the master controller, check this for the port on the master controller that connects to the local controller.

Configuring Local Controller Settings

Many controller settings are unique to that device and therefore are not replicated from a master controller to a local controller. The following settings must be manually configured on a local controller that synchronizes with the master controller.

- Time zone and daylight savings time settings
- VPN pools for remote APs and other VPN clients.
- Controller and IP interfaces. (Note that these values may need to be set *before* synchronization with the master so the synchronization can properly complete.)
- IP routing and spanning-tree configurations
- Remote AP whitelist and local-user database values



By default, the local controllers forward the authentication requests for the RAP whitelist and the local user database to the master controller. Therefore, this data does not have to be manually replicated *unless* the default behavior has been altered. The user table is NOT synchronized, so if an AP fails over to a master from a local or vice versa, that AP will have to re-authenticate.

- DHCP pools
- NAT pools
- SNMP, NTP, and syslog settings
- Hostnames, DNS and SMTP servers
- ACLs applied to ports

- Certificates
- RADIUS client details and RADIUS source interfaces
- Stateful firewall settings
- Customized captive portal pages and images, and the captive portal redirect address.



If you want to configure GRE tunnel between master and local controllers, you should use controller-IPs as tunnel endpoints.

Configuring APs

APs download their configurations from a master controller. However, an AP or AP group can tunnel client traffic to a local controller. To specify the controller to which an AP or AP group tunnels client traffic, you configure the LMS IP in the AP system profile on the master controller.

Configuration changes take effect only after you reboot the affected APs; this allows them to reassociate with the local controller. After rebooting, these APs appear to the new local controller as local APs.

Using the WebUI to configure the LMS IP

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
   - If you select AP Group, click Edit for the AP group name for which you want to configure the LMS IP.
   - If you select AP Specific, select the name of the AP for which you want to configure the LMS IP.
2. Under the Profiles section, select AP to display the AP profiles.
3. Select the AP system profile you want to modify.
4. Enter the controller IP address in the LMS IP field.
5. Click **Apply**.

Using the CLI to configure the LMS IP

```
ap system-profile <profile>
   lms-ip <ipaddr>

ap-group <group>
   ap-system-profile <profile>

ap-name <name>
   ap-system-profile <profile>
```

Moving to a Multi-Controller Environment

For a single WLAN configuration, the master controller is the controller which controls the RF and security settings of the WLAN. Additional controllers to the same WLAN serve as local switches to the master controller. The local controller operates independently of the master controller and depends on the master controller only for its security and RF settings. You configure the layer-2 and layer-3 settings on the local controller independent of the master controller. The local controller needs to have connectivity to the master controller at all times to ensure that any changes on the master are propagated to the local controller.

Some of the common reasons to move from a single to a multi-controller-environment include:

- Scaling to include a larger coverage area
- Setting up remote Access Points (APs)
- Network setup requires APs to be redistributed from a single controller to multiple controllers

You can use a preshared key (PSK) or a certificate to create IPSec tunnels between a master and backup master controllers and between master and local controllers. These inter-controller IPSec tunnels carry management traffic such as mobility, configuration, and master-local information.



An inter-controller IPSec tunnel can be used to route data between networks attached to the controllers if you have installed PEFV licenses in the controllers. To route traffic, configure a static route on each controller specifying the destination network and the name of the IPSec tunnel.

There is a default PSK to allow inter-controller communications, however, for security you need to configure a unique PSK for each controller pair. You can use either the WebUI or CLI to configure a 6-64 character PSK on master and local controllers. To configure a unique PSK for each controller pair, you must configure the master controller with the IP address of the local and the PSK, and configure the local controller with the IP address of the master and the PSK.

You can configure a global PSK for all master-local communications, although this is not recommended for networks with more than two controllers. On the master controller, use **0.0.0.0** for the IP address of the local. On the local controller, configure the IP address of the master and the PSK.

The local controller can be located behind a NAT device or over the Internet. On the local controller, when you specify the IP address of the master controller, use the public IP address for the master.

If your master and local controllers use a pre-shared key for authentication, the IPsec tunnel will be created using IKEv1. If they use a factory-installed or custom certificate, they will use IKEv2 to create the IPsec tunnel. Controllers using IKEv2 and custom-installed certificates can optionally use Suite-B encryption for IPsec encryption. For details and requirements for Suite-B encryption, see Suite-B Cryptography on page 415.

Configuring a Preshared Key

Leaving the PSK set to the default value exposes the IPSec channel to serious risk, therefore you should always configure a unique PSK for each controller pair.

Sharing the same PSK between more than two controllers increases the likelihood of compromise. If one controller is compromised, all controllers are compromised. Therefore, best security practices include configuring a unique PSK for each controller pair



Do not use the default global PSK on a master or stand-alone controller. If you have a multi-controller network then configure the local controllers to match the new IPSec PSK key on the master controller.

Weak keys are susceptible to offline dictionary attacks, meaning that a hostile eavesdropper can capture a few packets during connection setup and derive the PSK, thus compromising the connection. Therefore the PSK selection process should be the same process as selecting a strong passphrase:

- the PSK should be at least ten characters in length
- the PSK should not be a dictionary word
- the PSK should combine characters from at least three of the following four groups:
  - lowercase characters
  - uppercase characters
  - numbers
  - punctuation or special characters, such as !~‘@#$%^&*()_-+=\|//.[]{}

The following sections describe how to configure a PSK using the WebUI or CLI.

Using the WebUI to configure a Local Controller PSK

1. Navigate to the **Configuration > Network > Controller > System Settings** page.

2. The procedure to configure a local PSK varies, depending upon whether it is configured using a local controller or a master controller.
   - On a local controller, enter the IPSec key in the **IPSec Key (IKE PSK)** and **Retype IPSec Key (IKE PSK)** fields.
   - On a master controller, click **New** under **Local Controller IPSec Keys**. then enter the local controller IP address and then enter and retype the IPSec key. Click **Add**.
3. Click **Apply**.

Using the WebUI to configure a Master Controller PSK

Use the procedure below to configures the IP address and preshared key for the master controller.

1. Navigate to the **Configuration > Network > Controller > System Settings** page.
2. In the **IPSEC Key (IKE PSK)** field, enter the IPSec key. Reenter this key in the **Retype IPSEC Key (IKE PSK)** field.
3. (Optional) In the **FQDN** field, enter a fully qualified domain name used in IKE.
4. (Optional) Click the **Source IP address field** and select the VLAN ID of Vlan interface to initiate IKE. The controller IP address will be used if the VLAN is not specified.
5. Click **Apply**.

Using the CLI to configure a PSK

Master Controller

On the master controller you can configure a specific IPSec PSK for a local controller and use the localip 0.0.0.0 ipsec command:



You need to change the secret key to a non-default PSK key value even if you use a per-local controller PSK key configuration.

```
localip 0.0.0.0 ipsec <secret_key>
localip <ipaddr> ipsec <secret_key>
```

Local Controller

On the local controller the secret key (PSK) must match the master controller's PSK.

```
masterip <ipaddr> ipsec <secret_key> [fqdn <fqdn>][uplink][vlan <id>]
```

Configuring a Controller Certificate

The following sections describe how to use the command-line interface to select a factory-installed or custom certificate for secure inter-controller communication.

Using the CLI to configure a Local Controller Certificate

- Issue the following command on a master controller to configure the factory-installed certificate for secure communication between that master and a local controller.

  ```
  local-factory-cert local-mac <lmac>
  ```

  In this command, **<lmac>** is the MAC address of the local controller's factory-installed certificate.
- Issue the following command on a master controller to configure a custom certificate for secure communication between that master and a local controller.

  ```
  local-custom-cert local-mac <lmac> ca-cert <ca> server-cert <cert>

  suite-b <gcm-128 | gcm-256>
  ```

  In this command, **<lmac>** is the MAC address of the local controller's custom certificate.

Using the CLI to configure the Master Controller Certificate

Issue the following command on a local controller to configure the preshared key or certificate for the master controller.

```
masterip <ipaddr>
   ipsec <key> [interface uplink|{vlan <id>}] [fqdn <fqdn>]
   ipsec-custom-cert master-mac1 <mac1> [master-mac2 <mac2>] ca-cert <ca> server-cert <cert>
   [interface uplink|{vlan <id>}] [fqdn <fqdn>] [suite-b gcm-128|gcm-256]
   ipsec-factory-cert master-mac1 <mac1> [master-mac2 <mac2>] [interface uplink|{vlan <id>}]
   [fqdn <fqdn>]
```

Extreme Security (xSec) is a cryptographically secure, Layer-2 tunneling network protocol implemented over the 802.1x protocol. The xSec protocol can be used to secure Layer-2 traffic between the Dell controller and wired and wireless clients, or between Dell controllers.



xSec is an optional ArubaOS software module. You must purchase and install the license for the xSec software module on the controller.

Topics in this chapter include:

- Securing Client Traffic on page 863
- Securing Controller-to-Controller Communication on page 870
- Configuring the Odyssey Client on Client Machines on page 871

xSec encrypts an original Layer-2 data frame inside a Layer-2 xSec frame, the contents of which are defined by the protocol. xSec relies on 256-bit Advanced Encryption Standard (AES) encryption.

Upon 802.1x client authentication, xSec creates a tunnel between the client and the controller. The xSec frame sent over the air or wire between the user and the controller contains user and controller information, as well as original IP and MAC addresses, in encrypted form. All user information is secured using xSec. This concept is also extended to secure management information and data between two controllers on the same VLAN.

For xSec tunneling between a client and controller to work, a version of the Funk Odyssey client software that supports xSec needs to be installed on the client. It is possible to secure clients running Windows 2000 and XP operating systems using xSec and the Odyssey client software..



xSec is an optional licensed feature for Dell controllers. xSec is automatically enabled on the controller when you install the license. For information about the currently supported release for Funk Odyssey, please contact Juniper Networks.

xSec provides the following advantages:

- Advanced security as Layer-2 frames are encrypted and tunneled.
- Ease of implementation of advanced encryption in a heterogeneous environment. xSec is designed to support multiple operating systems and a wide range of network interface cards (NICs). All encryption and decryption on the client machine is performed by the Odyssey client while the NICs are configured with NULL encryption. This ensures that even older operating systems that cannot be upgraded to support WPA or WPA2 authentication can be secured using xSec and the Odyssey client.
- Compatible with TLS, TTLS and PEAP.
- Advanced authentication extended to wired clients allowing network managers to secure wired ports.

Securing Client Traffic

You can secure wireless or wired client traffic with xSec. On the client, install the Odyssey Client software. The xSec client must complete 802.1x authentication. to connect to the network. The client indicates the use of the xSec protocol during 802.1x exchanges with the controller. (Dell controllers support 802.1x for both wired and wireless clients.) Upon successful client authentication, an xSec tunnel is established between the controller and the client.

The authenticated client is placed into a configured VLAN, which determines the client's DHCP server, IP address, and Layer-2 connection. For wireless xSec clients, the VLAN is the user VLAN configured for the WLAN. For wired xSec clients and wireless xSec clients that connect to the controller through a non-Dell AP, the VLAN is a designated xSec VLAN. The VLAN can also be derived from configured RADIUS server-derivation rules or from Vendor-Specific Attributes (VSAs). Once an xSec tunnel is established, a DHCP server assigns the xSec client an IP address from the address pool on the VLAN to which the client is assigned. All traffic between the client and the controller is then encrypted.

The following sections describe how to configure xSec on the controller for wireless and wired clients.

Securing Wireless Clients

The following are the basic steps for configuring the controller for xSec wireless clients:

1. Configure the user VLAN to which the authenticated clients will be assigned. See Network Configuration Parameters on page 148 for more information.
2. Configure the user role for the authenticated xSec clients. See Roles and Policies on page 364for information.
3. Configure the server group that will be used to authenticate clients using 802.1x. See Authentication Servers on page 225 for more information
4. Configure the AAA profile to specify the 802.1x default user role. Specify the 802.1x authentication server group.

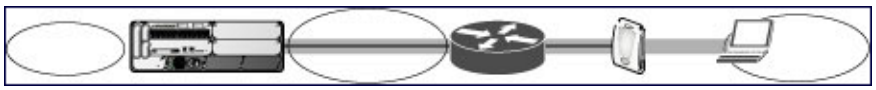


You can configure the 802.1x authentication profile if necessary. See 802.1X Authentication on page 253 for more information.

5. Configure the virtual AP profile for the WLAN. Specify the previously-configured user VLAN. Only xSec clients will be allowed to connect to the WLAN and non-xSec connections are dropped.
   a. Specify the previously-configured AAA profile.
   b. Configure the SSID profile with xSec as the authentication.
6. Install and set up the Odyssey Client on the wireless client.

Figure 178 is an example network where a wireless xSec client is assigned to the user VLAN 20 and the user role "employee" upon successful 802.1x authentication. VLAN 1 includes the port on the controller that connects to the wired network on which the AP is installed. (APs can connect to the controller across either a Layer-2 or Layer-3 network.)

Figure 178 *Wireless xSec Client Example*



The following sections describe how to use the WebUI or CLI to configure the AAA profile and virtual AP profile for this example. Other chapters in this manual describe the configuration of the user role, VLAN, authentication servers and server group, and 802.1x authentication profile.

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > AAA Profiles** page.
   a. To create a new AAA profile, click **Add** in the AAA Profiles Summary.
   b. Enter a name for the profile (for example, **xsec-wireless**), and click **Add**.
   c. To configure the AAA profile, click on the newly-created profile name.
   d. For 802.1x Authentication Default Role, select a configured user role (for example, **employee**).

e. Click **Apply**.

f. In the AAA Profile list, select 802.1x Authentication Profile under the AAA profile you configured. Select the applicable 802.1x authentication profile (for example, **xsec-wireless-dot1x**). Click **Apply**.

g. In the AAA Profile list, select 802.1x Authentication Server Group under the AAA profile you configured. Select the applicable server group (for example, **xsec-svrs**). Click **Apply**.

2. Navigate to the **Configuration > Wireless > AP Configuration** page. Select either the AP Group or AP Specific tab. Click **Edit** for the applicable AP group name or AP name.
3. Under Profiles, select Wireless LAN, then select Virtual AP.
4. To create a new virtual AP profile, select NEW from the Add a profile drop-down menu. Enter the name for the virtual AP profile (for example, **xsec-wireless**), and click **Add**.

   a. In the Profile Details entry for the new virtual AP profile, select the AAA profile you previously configured. A pop-up window displays the configured AAA profile parameters. Click **Apply** in the pop-up window.

   b. From the SSID profile drop-down menu, select NEW. A pop-up window allows you to configure the SSID profile.

   c. Enter the name for the SSID profile (for example, **xsec-wireless**).

   d. Enter the Network Name for the SSID (for example, **xsec-ap**).

   e. For Network Authentication, select **xSec**.

   f. Click **Apply** in the pop-up window.

   g. At the bottom of the Profile Details page, click **Apply**.
5. Click on the new virtual AP name in the Profiles list or in Profile Details to display configuration parameters.

   a. Make sure Virtual AP enable is selected.

   b. For VLAN, enter the ID of the VLAN in which authenticated xSec clients are placed (for example, **20**).

   c. Click **Apply**.

In the CLI

```
aaa profile xsec-wireless
   authentication-dot1x xsec-wireless-dot1x
   d>ot1x-default-role employee
   d>ot1x-server-group xsec-svrs
wlan ssid-profile xsec-wireless
   essid xsec-ap
   opmode xSec
wlan virtual-ap xsec-wireless
   vlan 20
   aaa-profile xsec-wireless
   ssid-profile xsec-wireless
```

Securing Wired Clients

The following are the basic steps for configuring the controller for xSec wired clients:

1. Configure the VLAN to which the authenticated clients will be assigned. See Network Configuration Parameters on page 148 for information.

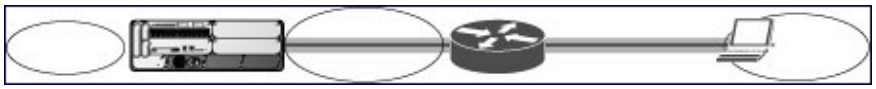
   This VLAN must have an IP interface, and is a different VLAN from the port's "native" VLAN that provides connectivity to the network.
2. Configure the user role for the authenticated xSec clients. See Roles and Policies on page 364 for information.
3. Configure the server group that will be used to authenticate clients using 802.1x. See Authentication Servers on page 225 for more information.

4. Configure the controller port to which the wired clients) are connected. Specify the VLAN to which the authenticated xSec clients are assigned.

   For firewall rules to be enforced after client authentication, the port must be configured as untrusted.
5. Configure the AAA profile to specify the 802.1x default user role and the 802.1x authentication server group.
6. Configure the wired authentication profile to use the AAA profile.
7. Install and set up the Odyssey Client on the wireless client.

Figure 179 is an example network where a wired xSec client is assigned to the VLAN 20 and the user role "employee" upon successful 802.1x authentication. Traffic between the controller and the xSec client is encrypted.

Figure 179 *Wired xSec Client Example*



The VLAN to which you assign an xSec client must be a different VLAN from the VLAN that contains the controller port to which the wired xSec client or AP is connected.

The following sections describe how to use the WebUI or CLI to configure the controller port to which the wired client is connected, the AAA profile, and the wired authentication profile for this example. Other chapters in this manual describe the configuration of the user role, VLAN, authentication servers and server group, and 802.1x authentication profile.

In the WebUI

1. Navigate to the **Configuration > Networks > Ports** page to configure the port to which the wired client(s) are connected.
   a. Click the port that you want to configure.
   b. Make sure the Enable Port checkbox is selected.
   c. For Enter VLAN(s), select the native VLAN on the port to ensure Layer-2 connectivity to the network. In Figure 179, this is VLAN 1.
   d. For xSec VLAN, select the VLAN to which authenticated users are assigned from the drop-down menu. In Figure 179, this is VLAN 20.
   e. Click **Apply**.
2. Navigate to the **Configuration > Security > Authentication > AAA Profiles** page to configure the AAA profile.
   a. To create a new AAA profile, click **Add**.
   b. Enter a name for the profile (for example, **xsec-wired**), and click **Add**.
   c. To configure the AAA profile, click on the newly-created profile name.
   d. For 802.1x Authentication Default Role, select a configured user role (for example, **employee**).
   e. Click **Apply**.
   f. In the AAA Profile list, select 802.1x Authentication Profile under the AAA profile you configured. Select the applicable 802.1x authentication profile (for example, **xsec-wired-dot1x**). Click **Apply**.
   g. In the AAA Profile list, select 802.1x Authentication Server Group under the AAA profile you configured. Select the applicable server group (for example, **xsec-svrs**). Click **Apply**.
3. Navigate to the **Configuration > Advanced Services > Wired Access** page.
   a. Under Wired Access AAA Profile, select the AAA profile you just configured.
   b. Click **Apply**.

In the CLI

```
interface fastethernet|gigabitethernet slot/port
   switchport access vlan 1
   xsec vlan 20
aaa profile xsec-wired
   authentication-dot1x xsec-wired-dot1x
   d>ot1x-default-role employee
   d>ot1x-server-group xsec-svrs
aaa authentication wired
   profile xsec-wired
```

Securing Wireless Clients Through Non-Dell APs

If xSec clients are connecting through a non-Dell AP, you need to configure the controller port to which the AP is connected. The AP must be configured for no (opensystem) authentication.

The following are the basic steps for configuring the controller for xSec wireless clients connecting through a non-Dell AP:

1. Configure the VLAN to which the authenticated clients will be assigned. See Network Configuration Parameters on page 148for information.

   This VLAN must have an IP interface, and is a different VLAN from the port's "native" VLAN that provides connectivity to the network.
2. Configure the user role for the authenticated xSec clients. See Roles and Policies on page 364 for information.
3. Configure the server group that will be used to authenticate clients using 802.1x. See Authentication Servers on page 225 for more information.
4. Configure the controller port that connects to the wired network on which the non-Dell AP is installed. Specify the VLAN to which the authenticated xSec clients are assigned.

   The ingress and egress ports for xSec client traffic must be different physical ports on the controller.
5. Configure the AAA profile to specify the 802.1x default user role and the 802.1x authentication server group.
6. Configure the wired authentication profile to use the AAA profile.
7. Install and set up the Odyssey Client on the wireless client.

The following sections describe how to use the WebUI or CLI to configure the controller port and AAA and wired authentication profiles for wireless clients connecting with non-Dell APs. Other chapters in this manual describe the configuration of the user role, VLAN, authentication servers and server group, and 802.1x authentication profile.

In the WebUI

1. Navigate to the **Configuration > Networks > Ports** page to configure the port to which the wireless xSec client(s) are connected.
   a. Click the port that you want to configure.
   b. Make sure the Enable Port checkbox is selected.
   c. For Enter VLAN(s), select the native VLAN (for example, **VLAN 1**) on the port to ensure Layer-2 connectivity to the network.
   d. For xSec VLAN, select the VLAN to which authenticated users are assigned from the drop-down menu (for example, **VLAN 20**)
   e. Click **Apply**.
2. Navigate to the **Configuration > Security > Authentication > AAA Profiles** page to configure the AAA profile.

a. To create a new AAA profile, click **Add**.
b. Enter a name for the profile (for example, **xsec-3party**), and click **Add**.
c. To configure the AAA profile, click on the newly-created profile name.
d. For 802.1x Authentication Default Role, select a configured user role (for example, **employee**).
e. Click **Apply**.
f. In the AAA Profile list, select 802.1x Authentication Profile under the AAA profile you configured. Select the applicable 802.1x authentication profile (for example, **xsec-NonDell-dot1x**). Click **Apply**.
g. In the AAA Profile list, select 802.1x Authentication Server Group under the AAA profile you configured. Select the applicable server group (for example, **xsec-svrs**). Click **Apply**.

3. Navigate to the **Configuration > Advanced Services > Wired Access** page.
   a. Under Wired Access AAA Profile, select the AAA profile you just configured.
   b. Click **Apply**.

In the CLI

```
interface fastethernet|gigabitethernet slot/port
   switchport access vlan 1
   xsec vlan 20
aaa profile xsec-wired
   authentication-dot1x xsec-NonDell-dot1x
   d>ot1x-default-role employee
   d>ot1x-server-group xsec-svrs
aaa authentication wired
   profile xsec-wired
```

Securing Clients on an AP Wired Port

APs with multiple wired Ethernet ports include an wired port profile that can enable or disable the wired port, define an AAA profile for wired port devices, and associate the port with an Ethernet link profile that defines its speed and duplex values.

In the WebUI

The procedure to create a new Ethernet port configuration profile depends upon whether or not you want to immediately associate that profile to a specific port on an AP.

1. To configure a new Ethernet port configuration profile without assigning it to a specific port:
   a. Navigate to the **Configuration > All Profiles** page.
   b. Expand the **AP** menu and select **AP Wired Port profile**.
   c. In the **Profile Details** window, enter a name for the new profile, then click **Add**.

   -or-

   To create a new Ethernet port configuration profile for a specific port on an AP or group of APs:
   a. Navigate to the **Configuration > Wireless > AP Configuration** page.
   b. Select either the **AP Group** or **AP Specific** tab. Click the **Edit** button by name of the AP group or individual AP you want to configure.
   c. In the Profiles list, expand the **AP**profile menu and select the **Ethernet Interface Port Configuration** profile for the Ethernet port number you want to configure.
   d. In the **Profile Details** window, click the **Ethernet interface port configuration** drop-down list and select **New**.
2. **Configure** the **Ethernet Interface Port/ Wired AP Port Configuration** profile parameters described in Table 186.

Table 186: *Ethernet Interface Port/ Wired AP Port Configuration Parameters*

Parameter	Description
`Shut Down`	Disable the wired AP port.
`Remote AP Backup`	If enabled, the port of Remote-AP is up for the local connectivity and troubleshooting when the controller is not reachable and no firewall policies will be applied. If disabled, the port would be up for the bridge mode when controllerr is not reachable, and retains the previous bridge wired port configuration (if the configuration is applied and persistent). For split and tunnel modes, the ports would be shutdown when the controller is not reachable.
Time to wait for authentication to succeed	Authentication timeout value, in seconds, for devices connecting the AP's wired port. The supported range is 1-65535 seconds, and the default value is 20 seconds.
Spanning Tree	Select this checkbox to enable the Spanning Tree protocol.

3. Each Ethernet Interface Port/AP Wired Port Configuration profile is automatically associated to the wired AP profile **Default**. To assign a new wired AP profile to the AP wired port:
   a. Click the wired AP profile directly under the Ethernet port profile you are editing.
   b. In the **Profile Details** window, click the **Wired AP Profile** drop-down list and select a new Wired AP profile.
4. A new AP wired profile is automatically associated to the Ethernet Interface Link profile **Default**. To assign a new Ethernet Interface Link profile to the AP wired port:
   a. Click the Ethernet Interface Link profile directly under the Ethernet port profile you are editing.
   b. In the **Profile Details** window, click the **Ethernet Interface Link** drop-down list and select a new Ethernet Interface Link profile.
5. By default, there is no AAA profile associated with an AP wired port profile. To assign an AAA profile to the AP wired port:
   a. Click the AAA profile directly under the Ethernet port profile you are editing.
   b. In the **Profile Details** window, click the **AAA Profile** drop-down list and select an AAA profile.
6. Click **Apply** to save your settings.

In the CLI

To create a new Ethernet Port/Wired AP Port profile, access the command-line interface in Config mode and issue the following command.

```
ap wired-port-profile <profile>
   aaa-profile <profile>
   authentication-timeout <seconds>
   enet-link-profile <profile>
   rap-backup
   shutdown
   wired-ap-profile <profile>
```

To associate an existing Ethernet Port/Wired AP Port profile to a specific interface on an AP or group of APs, access the command-line interface in Config mode and issue the following command.

```
ap-group <group>
   enet0-port-profile <profile>
   enet1-port-profile <profile>
```

```
enet2-port-profile <profile>
enet3-port-profile <profile>
enet4-port-profile <profile>
```

Enabling or Disabling the Spanning Tree Parameter in AP Wired Port Profile

You can enable or disable the Spanning Tree parameter in WebUI and CLI.

Using the WebUI

The following procedure configures the Spanning Tree parameter in AP Wired Port profile:

1. Navigate to the **Configuration > Advanced Services > All Profiles** page.
2. Under **AP > AP Wired Port** on the **Profiles** pane, select the profile name.
3. On the **Profile Details** pane, select the **Spanning Tree** check box.
4. Click **Apply**.

Using the CLI

The following example enables spanning tree in default ap-wired port profile, using the CLI command:

```
(host) (config) #ap wired-port-profile default
(host) (AP wired port profile "default") #spanning-tree
```

The following example displays the spanning tree information of an AP, using the CLI command:

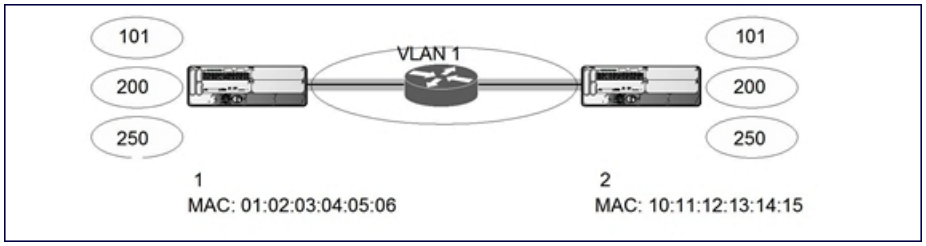
```
(host) (config) #show ap debug spanning-tree ap-name <ap-name>
```

Securing Controller-to-Controller Communication

xSec can be used to secure data and control traffic passed between two controllers. The only requirement is that both controllers be members of the same VLAN. To establish a point-to-point tunnel between the two controllers, you need to configure the following for the connecting ports on each controller:

- The MAC address of the xSec tunnel termination point. This would be the MAC address of the “other” controller.
- A 16-byte shared key used to authenticate the controllers to each other. You must configure the same shared key on both controllers.
- The VLAN IDs for the VLANs that will extend across both the controllers via the xSec. Figure 180 shows an example network where two controllers are connected to the same VLAN, VLAN 1. On controller 1, you configure the MAC address of controller 2 for the xSec tunnel termination point. On controller 2, you configure the MAC address of controller 1 for the xSec tunnel termination point. On both controllers, you configure the same 16-byte shared key and the IDs for the VLANs which are allowed to pass through the xSec tunnel.

Figure 180 *Controller-to-Controller xSec Example*



Configuring Controllers for xSec

The following sections describe how to use the WebUI or CLI to configure the port that connects to the wired network on which the other controller is installed. Other chapters in this manual describe the configuration of VLANs.

In the WebUI

1. On each controller, navigate to the **Configuration > Network > Port** page.
2. Click on the port to be configured.
3. Select the VLAN from the drop-down list.
4. Configure the xSec point-to-point settings:
   a. Enter the MAC address of the tunnel termination point (the "other" controller's MAC address).
   b. Enter the key (for example, 1234567898765432) used by xSec to establish the tunnel between the controllers.
   c. Select the VLANs that would be allowed across the point-to-point connection from the Allowed VLANs drop-down menu, and click the <-- button.
5. Click **Apply**.

In the CLI

For Controller 1:

```
interface gigabitethernet|fastethernet slot/port
   vlan 1
   xsec point-to-point 10:11:12:13:14:15 1234567898765432 allowed vlan 101,200,250
```

For Controller 2:

```
interface gigabitethernet|fastethernet slot/port
   vlan 1
   xsec point-to-point 01:02:03:04:05:06 1234567898765432 allowed vlan 101,200,250
```

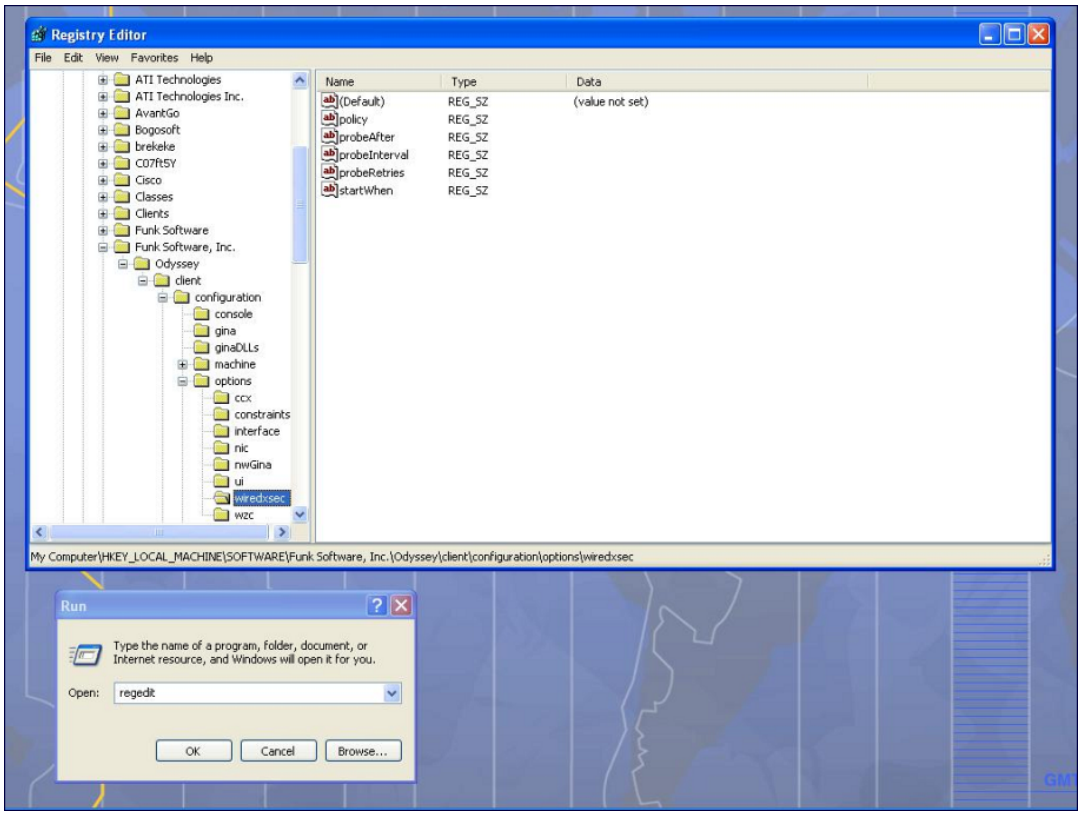
Configuring the Odyssey Client on Client Machines

You can obtain the Odyssey Client from Juniper Networks. For information on Odyssey Client versions, contact Juniper Networks Support.

Installing the Odyssey Client

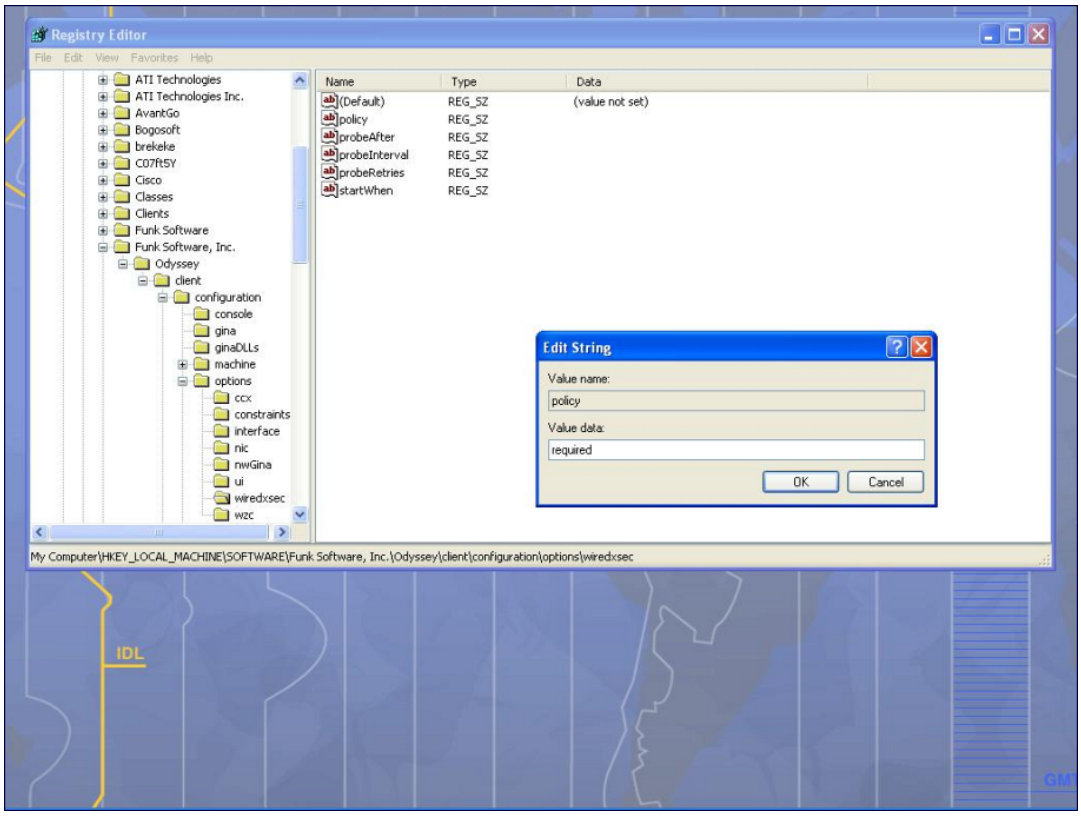
1. Unzip and install the Odyssey client on the client laptop.
2. For wired xSec, to use the Odyssey client to control the wired port, modify the registry:
   a. On the windows machine, click **Start** and select **Run**.
   b. Type **`regedit`** in the dialog box and click **OK**.
   c. Navigate down the tree to
   `HKEY_LOCAL_MACHINE\SOFTWARE\Funk Software, Inc.\odyssey\client\configuration\options\wiredxsec`.

Figure 181 *The regedit Window*



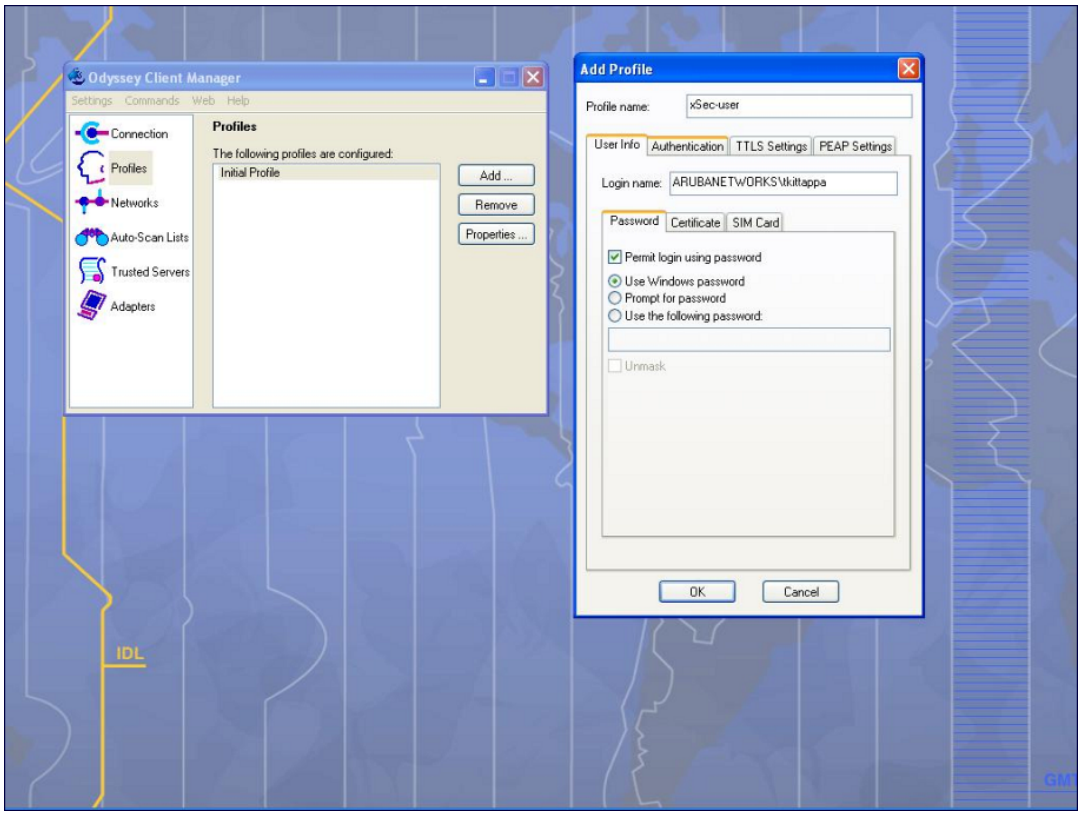
d. Select “policy” from the registry values and right click on it. Select **Modify** to modify the contents of policy. Set the value in the resulting window to **required**.

Figure 182 *Modifying a regedit Policy*



3. Open the Funk Odyssey Client. Click the **Profile** tab in the client window. This allows the user to create the user profile for 802.1x authentication.

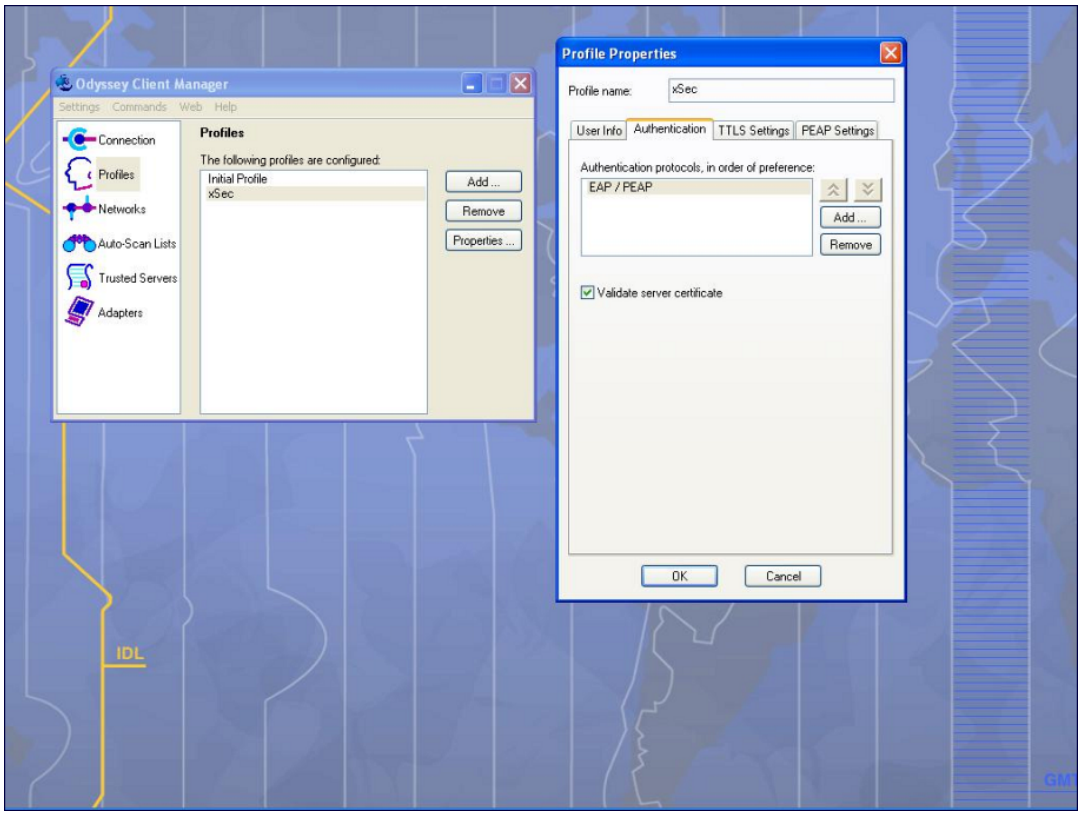
Figure 183 *The Funk Odyssey Client Profile*



a. In the login name dialog box, enter the login name used for 802.1x authentication. For the password, the client could use the WINDOWS password or use the configured password based on the selection made.

b. Click the certificate tab and enter the certificate information required. This example shows the PEAP settings.

Figure 184 *Certificate Information*



c. Click the **Authentication** tab. In the resultant window, click the **Add** tab and select **EAP/PEAP**. Move this option to the top of the list if PEAP is the method chosen. If certification validation not required, uncheck the **Validate server certificates** setting.

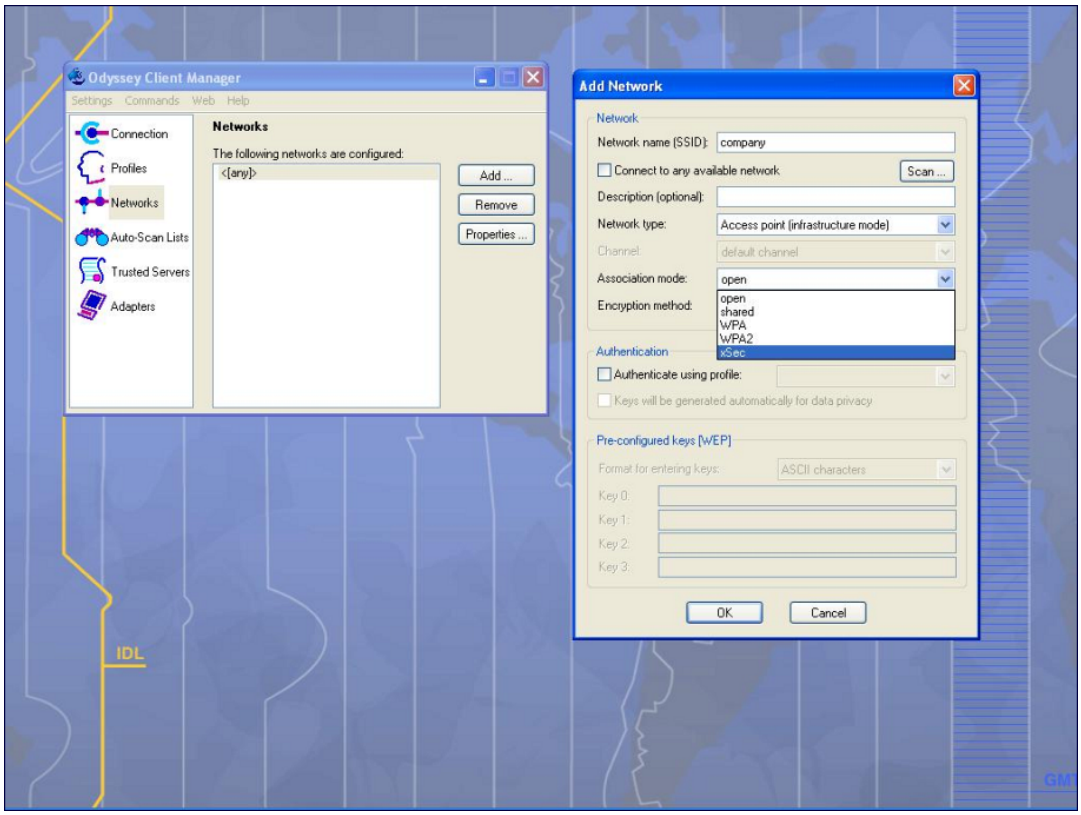
d. Click the **PEAP Settings** tab and select the EAP protocol supported.

e. Click **OK**.

f. To modify an existing profile, select the profile and then click the **Properties** tab.

4. Select the **Network** tab to configure the network for wireless client. For wired clients, skip this step.

Figure 185 *Network Profile*



a. Click the **Add** tab. Enter the SSID to which the client connects.
b. Set the Network type to **Infrastructure**.
c. Set the Association mode to **xSec**, AES encryption is automatically selected.
d. Under Authentication, select the **Authenticate using profile** checkbox.
e. From the pull down menu, select the profile used for 802.1x authentication. This would be one of the profiles configured in step 2.
f. Select the keys that will be generated automatically for data privacy.
g. Apply the configuration changes made by clicking on the **OK** tab.
h. To modify an existing profile, select the profile and then click the **Properties** tab.

5. Click the **Adapters** tab if the adapter used is not seen under the list of adapters pull down menu under connections.
   a. When using a wireless client, click the **Wireless** tab.
   b. Select the **Wireless adapters only** radio button. From the resulting list, select the adapter required from the list and click **OK**.
   c. For wired 802.1x clients, select the **Wired 802.1x** tab and select the **Wired adapters only** radio button. From the resulting list, select the adapter required from the list and click **OK**.
6. Establish the connection.
   a. Click the **Connection** tab.
   b. From the pull down menu, select the adapter required. If the adapter in use is not visible, add the adapter as explained in Step 5.
   c. Select the **Connect to network** checkbox and select the **Network** option from the pull down menu. To configure a new network, follow the instructions in Step 4.
   d. This will automatically start the connection process. To reconnect to the network, click **Reconnect**.

7. Click **Scan** to display the SSIDs seen by the NIC after a site survey.

This chapter outlines the steps required to configure voice and video services on the Dell controller for Voice over IP (VoIP) devices, including Session Initiation Protocol (SIP), Spectralink Voice Priority (SVP), H323, SCCP, Vocera, and Alcatel NOE phones, clients running Microsoft Lync Server, and Apple devices running the Facetime application. Because video and voice applications are more vulnerable to delay and jitter, the network infrastructure must be able to prioritize video and voice traffic over data traffic.

This chapter includes the following topics:

- Voice and Video License Requirements on page 878
- Configuring Voice and Video on page 878
- Working with QoS for Voice and Video on page 891
- Unified Communication and Collaboration on page 899
- Understanding Extended Voice and Video Features on page 915
- Advanced Voice Troubleshooting on page 933

Voice and Video License Requirements

The voice and video services require PEFNG licenses on the controller. For complete details on the required licenses, see Software Licenses on page 130.

Configuring Voice and Video

This section describes the steps required to set up and configure voice features on aDell controller:

1. Set up net services
2. Configure roles
3. Configure firewall settings for voice and video ALGs
4. Configure other parameters depending on the need and environment



Assigning voice traffic to the high priority queue is recommended when deploying voice over WLAN networks.

Voice ALG and Network Address Translation

Voice ALGs do not support Network Address Translation (NAT). This is due to the way NAT functions and the way IP addresses are embedded in the signaling messages. In a typical customer deployment, a call server is deployed within an internal network which eliminates the need for NAT.

In short, voice ALGs should not be enabled when voice clients are behind a NAT.

Setting up Net Services

You can either use the default net services and ports or you can create or modify net services.

Using Default Net Services

The following table lists the default net services and their ports:

Table 187: *Default Voice Net Services and Ports*

Net Service Name	Protocol	Port	ALG
svc-sccp	TCP	2000	SCCP
svc-sip-tcp	TCP	5060	SIP
svc-sip-udp	—	—	SIP
svc-sips	—	—	SIP
svc-noe	UDP	32512	NOE
svc-h323-udp	UDP	1718, 1719	H.323
svc-h323-tcp	TCP	1720	H.323
svc-vocera	—	—	VOCERA
svc-svp	—	None	SVP

Creating Custom Net Services

You can use CLI to create or modify net services. In the config mode on the controller enter:

```
(host) (config)# netservice [service name] [protocol] [port] [alg]
```

To create an *svc-noe* service on UDP port 32522, enter:

```
(host) (config)# netservice svc-noe udp 32522 alg noe
```

Configuring User Roles

In the user-centric network, the user role of a wireless client determines its privileges and the type of traffic that it can send or receive in the wireless network. You can configure roles for clients that use mostly data traffic, such as laptops, and roles for clients that use mostly voice traffic, such as VoIP phones. Although there are different ways for a client to derive a user role, in most cases the clients using data traffic are assigned a role after they are authenticated through a method such as 802.1x, VPN, or captive portal. The user role for VoIP phones is derived from the OUI of their MAC addresses or the SSID to which they associate. Refer to Roles and Policies on page 364 for details on how to create and configure a user role.

This section describes how to configure voice user roles with the required privileges and priorities. Dell controller provides default user roles for all voice services. You can do one of the following:

- Use default user roles
- Create or modify user roles
- Use user-derivation roles

Using the Default User Role

The controller is configured with the default voice role. This role has the following settings:

- No limit on upload or download bandwidth
- Default L2TP and PPTP pool
- Maximum sessions: 65535

The following ACLs are associated with the default voice role:

- SIP-ACL
- NOE-ACL
- SVP-ACL
- VOCERA-ACL
- SKINNY-ACL
- H323-ACL
- DHCP-ACL
- TFTP-ACL
- DNS-ACL
- ICMP-ACL

For more details on the default voice role, enter the following command in the config mode on your controller:

```
(host) (config) #show rights voice
```

Creating or Modifying Voice User Roles

You can create roles for NOE, SIP, SVP, Vocera, SCCP, and H.323 ALGs. Use the WebUI or CLI to configure user roles for any of the ALGs.

Using the WebUI to configure user roles

1. Navigate to the **Configuration > Security > Access Control** page.
2. Select the **Policies** tab. Click **Add** to create a new policy.
3. For **Policy Name**, enter a name here.
4. For **Policy Type**, select **Session**.
5. Under **Rules**, click **Add**.
   a. For **IP Version**, select **IPv4**.
   a. For **Source**, select **any**.
   b. For **Destination**, select **any**.
   c. For **Service**, select service, then the correct voice or video ALG service. See Table 188 and Table 189 for service names for all ALGs:

Table 188: *Services for ALGs*

ALG	Service Name
NOE	svc-noe sip-noe-oxo
SIP	• svc-sips • svc-sip-tcp • svc-sip-udp
SVP	svc-svp
VOCERA	svc-vocera
SCCP	svc-sccp
H.323	• svc-h323-tcp • svc-h323-udp

ALG	Service Name
DHCP	svc-dhcp
TFTP	svc-tftp
ICMP	svc-icmp
DNS	svc-dns

Table 189: *Other Mandatory Services for the ALGs*

ACL	Service Name
DHCP	svc-dhcp
TFTP	svc-tftp
ICMP	svc-icmp
DNS	svc-dns

d. For **Action**, select **permit**.
e. For **Queue**, select **High**.
f. Click **Add**. Repeat steps 1 to 5e to add more ALG services.

6. Click **Apply**.
7. Select the **User Roles** tab. Click **Add** to add a user role.
   a. For **Role Name**, enter a name for the user role.
   b. Under **Firewall Policies**, click **Add**.
   c. Select the previously-configured policy name from the **Choose from Configured Policies** drop-down menu.
   d. Click **Done**.
   e. Under **Firewall Policies**, click **Add**.
   f. Select **control** from the **Choose from Configured Policies** drop-down menu.
   g. Click **Done**.
8. Click **Apply**.

Using the CLI to configure a user role

Use the following commands:

```
ip access-list session <policy-name>
   any any <service-name> permit queue high
   any any dhcp-acl permit queue high
   any any tftp-acl permit queue high
   any any dns-acl permit queue high
   any any icmp-acl permit queue high

user-role <role-name>
   session-acl <policy-name>
```

Replace the following strings:

- *policy-name* with a string that you want to identify the roles policy

- *role-name* with the name you want to identify the voice user role
- *service-name* with any of the service names from Table 187

Using the User-Derivation Roles

The user role can be derived from attributes from the client's association with an AP. For VoIP phones, you can configure the devices to be placed in their user role based on the SSID or the Organizational Unit Identifier (OUI) of the client's MAC address.



User-derivation rules are executed *before* the client is authenticated.

Using the WebUI to Derive the Role Based on SSID

1. Navigate to the **Configuration > Security > Authentication > User Rules** page.
2. Click **Add** to add a new set of derivation rules. Enter a name for the set of rules, and click **Add**. The name appears in the **User Rules Summary** list.
3. In the **User Rules Summary** list, select the name of the rule set to configure rules.
4. Click **Add** to add a rule. For **Set Type**, select **Role** from the drop-down menu.
5. For **Rule Type**, select **ESSID**.
6. For **Condition**, select **equals**.
7. For **Value**, enter the SSID used for the phones.
8. For **Roles**, select the user role you previously created.
9. Click **Add**.
10. Click **Apply**.

Using the CLI to Derive the Role Based on SSID

Use the following commands:

```
aaa derivation-rules user name
   set role condition essid equals ssid set-value role
```

Using the WebUI to Derive the Role Based on MAC OUI

1. Navigate to the **Configuration > Security > Authentication > User Rules** page.
2. Click **Add** to add a new set of derivation rules. Enter a name for the set of rules, and click **Add**. The name appears in the **User Rules Summary** list.
3. In the **User Rules Summary** list, select the name of the rule set to configure rules.
4. Click **Add** to add a rule. For **Set Type**, select **Role** from the drop-down menu.
5. For **Rule Type**, select **MAC Address**.
6. For **Condition**, select **contains**.
7. For **Value**, enter the first three octets (the OUI) of the MAC address of the phones (for example, the Spectralink OUI is 00:09:7a).
8. For **Roles**, select the user role you previously created.
9. Click **Add**.
10. Click **Apply**.

Using the CLI to Derive the Role Based on MAC OUI

Use the following commands:

```
aaa derivation-rules user name
   set role condition macaddr contains xx:xx:xx set-value role
```

Configuring Firewall Settings for Voice and Video ALGs

After configuring the user roles, you must configure the firewall settings for the voice and video Application-Level Gateways (ALGs) to pass traffic securely through the Dell devices.

You can use the WebUI or CLI to configure the firewall settings for the ALGs.

In the WebUI

1. Navigate to the **Configuration > Advanced Services > Stateful Firewall** page.
2. Enable the firewall settings for the ALGs:
   a. Select the **Stateful SIP Processing** check box for the SIP ALG.
   b. Select the **Stateful H.323 Processing** check box for the H.323 ALG.
   c. Select the **Stateful SCCP Processing** check box for the SCCP ALG.
   d. Select the **Stateful Vocera Processing** check box for the Vocera ALG.
   e. Select the **Stateful UA Processing** check box for the NOE ALG.

In the CLI

To enable the firewall settings for the SIP ALG:

```
(host) #configure terminal
(host) (config) #no firewall disable-stateful-sip-processing
```

To enable the firewall settings for the H.323 ALG:

```
(host) (config) #no firewall disable-stateful-h323-processing
```

To enable the firewall settings for the SCCP ALG:

```
(host) (config) #no firewall disable-stateful-sccp-processing
```

To enable the firewall settings for the Vocera ALG:

```
(host) (config) #no firewall disable-stateful-vocera-processing
```

To enable the firewall settings for the NOE ALG:

```
(host) (config) #no firewall disable-stateful-ua-processing
```

Additional Video Configurations

You can configure ArubaOS to reliably and efficiently stream video traffic over wireless LAN (WLAN). This new method allows you to stream video traffic reliably without much distortion. To ensure that video data is transmitted reliably, dynamic multicast optimization techniques are used.

Although the dynamic multicast optimization conversion generates more traffic, that traffic is buffered by the AP and delivered to the client when the client emerges from power-save mode.

Configuring Video over WLAN enhancements

To configure video over WLAN enhancements:

- Enable **WMM** on the SSID profile.
- Enable **IGMP** proxy or IGMP snooping.
- Configure an ACL to set a DSCP value same as the `wmm-vi-dscp` value in the SSID profile for prioritizing the multicast video traffic.
- Enable **dynamic multicast optimization** under **VAP profile**.
- Configure the dynamic multicast optimization threshold—The maximum number of high throughput stations in a multicast group. The optimization will stop if the number exceeds the threshold value.

- Enable **multicast rate optimization** to support higher data rate for multicast traffic in the absence of dynamic multicast optimization. Dynamic multicast optimization takes precedence over multicast rate optimization up to the configured threshold value.

NOTE

Configuring the **Video Multicast Rate Optimization** parameter overrides the configuration of **BC/MC Rate Optimization** parameter for VI-tagged multicast traffic. Multicast traffic that is not VI-tagged behaves the same with BC/MC as before. If multicast rate is not set, all traffic behaves the same.

- Enable **video aware scan** on **ARM profile**—This ensures that AP does not scan when a video stream is active.
- Optionally, you can configure and apply the **WMM bandwidth management profile**—The total bandwidth share should not exceed 100 percent.
- Enable multicast shaping to shape the sudden traffic from the source.

You can either use the WebUI or the CLI to configure the video over WLAN enhancements.

Prerequisites

- You will need the Policy Enforcement Firewall Next Generation (PEFNG) license to enable dynamic multicast optimization.
- This feature is available only on W-7200 Series, W-6000, W-3000 Series, and W-600 Series controller platforms.

In the WebUI

1. Enable IGMP proxy or IGMP snooping on the controller.

   To enable IGMP proxy:

   a. Navigate to the **Configuration > Network > IP** page. Under the **IGMP** settings, select the **Enable IGMP** checkbox.

   b. Select the **Proxy** checkbox and then the appropriate value from the **Interface** drop down menu.

   c. Click **Apply**.

Figure 186 *Enable IGMP Proxy*



   To enable IGMP snooping:

   a. Navigate to the **Configuration > Network > IP** page. Under the **IGMP** settings, select the **Enable IGMP** checkbox.

   b. Select the **Snooping** checkbox.

   c. Click **Apply**.

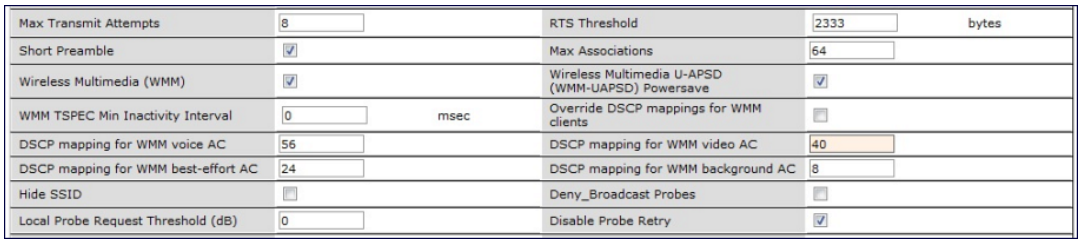
Figure 187 *Enable IGMP Snooping*



2. Enable wireless multimedia and set a DSCP value for video traffic.

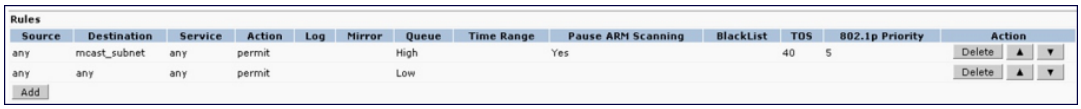
a. Navigate to the **Configuration > Advanced Services > All Profiles** page.
b. Under the **Profiles** column, expand **Wireless LAN > SSID Profile** and select the profile name. This example uses the *default* profile.
c. Click the **Advanced** tab and select the **Wireless Multimedia (WMM)** checkbox.
d. Enter the DSCP value (integer number) in the **DSCP mapping for WMM video AC field** and click **Apply**.

Figure 188 *Enable Wireless Multimedia and Set DSCP Value*



3. Create an ACL on the controller with the values equivalent to the DSCP mappings to prioritize the video traffic.
   a. Navigate to the **Configuration > Security > Access Control** page and click the **Policies** tab.
   b. Click **Add** to create a new policy.
   c. Enter the appropriate values under **Rules** to match the DSCP mapping values.

Figure 189 *Set ACL to Prioritize Video Traffic*

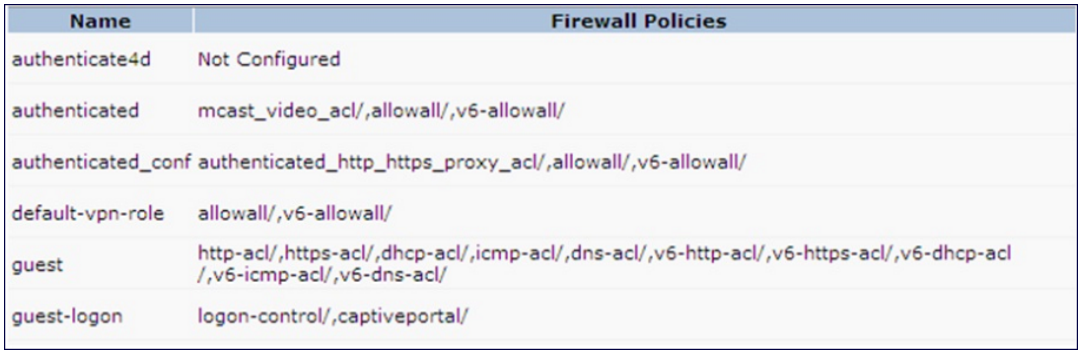


You can also add this ACL to any user role or port.

To apply the ACL to a user role:

a. Navigate to the **Configuration > Security > Access Control** page and click the **User Roles** tab.
b. Edit the user role and click **Add** under Firewall Policies.
c. Select the ACL from the **Choose From Configured Policies** drop-down menu and click the **Done** button.
d. Click **Apply** to save the configurations.

Figure 190 *Apply ACL to User Role*



To apply the ACL to a port:

a. Navigate to the **Configuration > Network > Port** page and select the upstream port.
b. Under the **VLAN Firewall Policy** drop-down menu, select the ACL.
c. Click **Apply** .

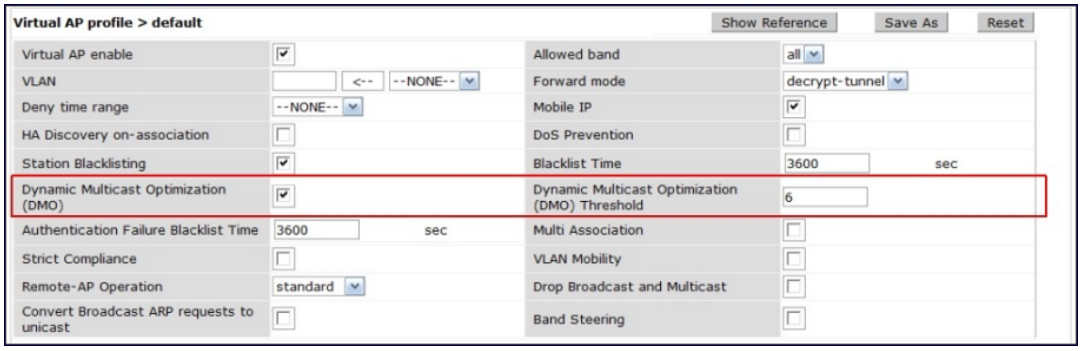
Figure 191 *Apply ACL to Port*



4. Configure dynamic multicast optimization for video traffic on a virtual AP profile.

   Under the **Profiles** column, expand **Wireless LAN > Virtual AP Profile** and select the profile name. This example uses the *default* profile. In the **Profile Details** section, select the **Dynamic Multicast Optimization (DMO)** option and enter the threshold value.

Figure 192 *Enabling Dynamic Multicast Optimization for Video and Set Threshold*

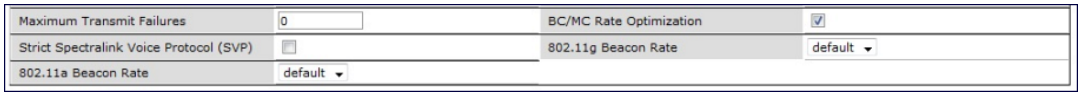


5. Configure multicast rate optimization for the video traffic.
   a. Navigate to the **Configuration > Advanced Services > All Profiles** page.
   b. Under the **Profiles** column, expand **Wireless LAN > SSID Profile** and select the profile name.
   c. Click the **Advanced** tab and select the **BC/MC Rate Optimization** checkbox.
   d. Select an option from **Video Multicast Rate Optimization** drop-down menu
   e. Click **Apply**.



Configuring the **Video Multicast Rate Optimization** parameter overrides the configuration of **BC/MC Rate Optimization** parameter for VI-tagged multicast traffic. Multicast traffic that is not VI-tagged behaves the same with BC/MC as before. If multicast rate is not set, all traffic behaves the same.

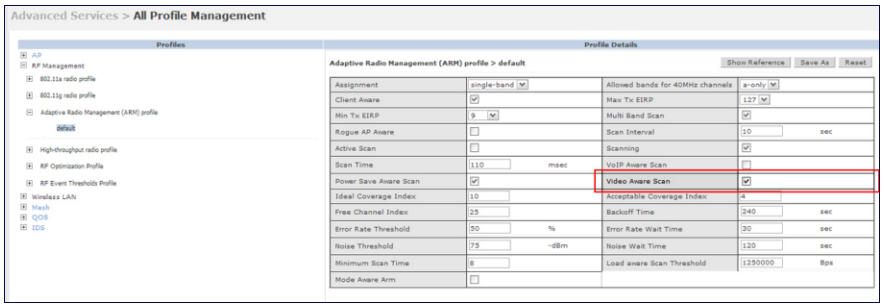
Figure 193 *Enable Multicast Rate Optimization*



6. Configure ARM scanning for video traffic.

   Under the **Profiles** column, expand **RF Management > Adaptive Radio Management (ARM) Profile** and select the profile name. This example uses the *default* profile. Select the **Video Aware Scan** option and click **Apply**.

Figure 194 *Enabling Video Aware Scan*



7. Configure and apply bandwidth management profile:

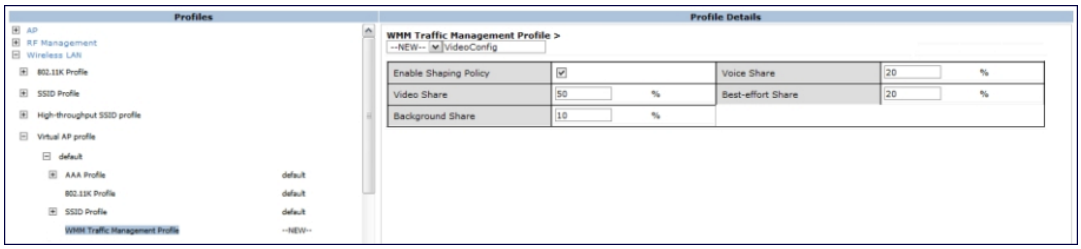
   Under the **Profiles** column, expand **Virtual AP >***[profile-name]* **> WMM Traffic Management Profile**. In the **Profile Details** section, select the profile name from the drop down list box. Select the **Enable Shaping Policy** option and enter the bandwidth share values. Click **Apply**.

   This step is optional.



Ensure that you configure the WMM traffic management profile to the virtual AP profile, if you have configured the virtual AP traffic management profile.

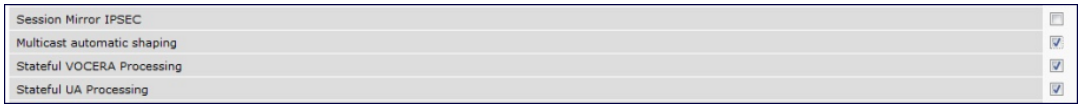
Figure 195 *Configuring bandwidth management*



After you configure the WMM bandwidth management profile, apply it to the virtual AP profile.

8. Enable multicast shaping on the firewall.
   a. Navigate to the **Configuration > Advanced Services > Stateful Firewall** page.
   b. Click the **Global Setting** tab and select the **Multicast automatic shaping** checkbox.
   c. Click **Apply**.

Figure 196 *Enable Firewall Multicast Shaping*



In the CLI

1. Enable IGMP proxy or IGMP snooping on the controller.

   To enable IGMP proxy:

```
(host) (config) #interface vlan 1
(host) (config-subif)#ip igmp proxy gigabitethernet 1/3
```

   To enable IGMP snooping:

```
(host) (config) #interface vlan 1
(host) (config-subif)#ip igmp snooping
```

2. Enable wireless multimedia and set a DSCP value for video traffic:

```
(host) (config)#wlan ssid-profile default
(host) (ssid-profile "default")#wmm
(host) (ssid-profile "default")#wmm-vi-dscp <value>
```

Example:

```
(host) (ssid-profile "default")#wmm-vi-dscp 40
(host) (SSID Profile "default") #show wlan ssid-profile default

SSID Profile "default"
----------------------
Parameter                                          Value
---------                                          -----
SSID enable                                        Enabled
ESSID                                              building1-ap
...
...
...
Wireless Multimedia (WMM)                          Enabled
Wireless Multimedia U-APSD (WMM-UAPSD) Powersave   Enabled
WMM TSPEC Min Inactivity Interval                  0 msec
Override DSCP mappings for WMM clients             Disabled
DSCP mapping for WMM voice AC                      56
DSCP mapping for WMM video AC                      40
...
...
...
```

Setting the DSCP value tags the content as video stream that the APs can recognize.

3. Create an ACL on the controller with the values equivalent to the DSCP mappings to prioritize the video traffic.

Example: The following ACL prioritizes the multicast traffic from the specified multicast group on the controller. You can also add this ACL to any user role or port:

```
(host) (config-sess-mcast_video_acl)#any network 224.0.0.0 255.0.0.0 any permit tos 40
queue high 802.1p 5
```

a. To add the ACL to a user role:

```
(host) (config) #user-role authenticated access-list session mcast_video_acl
```

This example uses the user role *authenticated*.

b. To add the ACL to a port:

```
(host) (config) #interface gigabitethernet 1/3
(host) (config-if)#ip access-group mcast_video_acl session
```

4. Configure dynamic multicast optimization for video traffic on a virtual AP profile:

```
(host) (config)#wlan virtual-ap default
(host) (Virtual AP Profile "default")#dynamic-mcast-optimization
(host) #show wlan virtual-ap default
Virtual AP profile "default"
---------------------------
Parameter                                  Value
---------                                  -----
Virtual AP enable                          Enabled
...
...
Blacklist Time                             3600 sec
Dynamic Multicast Optimization for Video                Enabled
Dynamic Multicast Optimization Threshold               6
...
```

```
...
```

5. Configure the dynamic multicast optimization threshold value:

```
(host) (config) #dynamic-mcast-optimization-thresh 6
(host) #(host} #show wlan virtual-ap default
Virtual AP profile "default"
---------------------------
Parameter                                   Value
---------                                   -----
Virtual AP enable                           Enabled
Allowed band                                all
...
...
...
Blacklist Time                              3600 sec
Dynamic Multicast Optimization for Video                Enabled
Dynamic Multicast Optimization Threshold               6
Authentication Failure Blacklist Time       3600 sec
...
...
...
```

6. Configure multicast rate optimization for video traffic:

```
(host) (config) #wlan ssid-profile default
(host) (SSID Profile "default") #mcast-rate-opt
(host) (SSID Profile "default") #show wlan ssid-profile default
SSID Profile "default"
----------------------
Parameter                                           Value
---------                                           -----
SSID enable                                         Enabled
ESSID                                               building1-ap
Encryption                                          opensystem
DTIM Interval                                       1 beacon periods
802.11a Basic Rates                                 6 12 24
...
...
...
EDCA Parameters Station profile                     N/A
EDCA Parameters AP profile                          N/A
BC/MC Rate Optimization                             Enabled
Strict Spectralink Voice Protocol (SVP)             Disabled
...
...
...
```

7. Configure ARM scanning for video traffic.

   In the default **RF ARM** profile, enable the **video aware scan** option. This prevents APs from scanning when a video traffic is active:

```
(host) (config) #rf arm-profile default
(host) (Adaptive Radio Management (ARM) profile "default") #video-aware-scan
(host) (Adaptive Radio Management (ARM) profile "default") #end
(host) #show rf arm-profile default

Adaptive Radio Management (ARM) profile "default"
-------------------------------------------------
Parameter                         Value
---------                         -----
Assignment                        single-band
Allowed bands for 40MHz channels  a-only
Client Aware                      Enabled
...
```

```
...
...
Scanning                          Enabled
Scan Time                         110 msec
VoIP Aware Scan                   Disabled
Power Save Aware Scan             Enabled
Video Aware Scan                  Enabled
...
...
...
Load aware Scan Threshold         1250000 Bps
Mode Aware Arm                    Disabled
```

8. Configure and apply a bandwidth management profile:

```
(host) (config)# wlan wmm-traffic-management-profile default
```



Ensure that you configure the WMM traffic management profile to the virtual AP profile, if you have configured the virtual AP traffic management profile.

a. Enable a bandwidth shaping policy so that the allocated bandwidth share is appropriately used:

```
(host) (WMM Traffic management profile "default") # enable-shaping
```

b. Set a bandwidth percentage for the following categories:

```
(host) (WMM Traffic management profile "default") # background 10
(host) (WMM Traffic management profile "default") # best-effort 20
(host) (WMM Traffic management profile "default") # video 50
(host) (WMM Traffic management profile "default") # voice 20
(host) (WMM Traffic management profile "default") # show wlan wmm-traffic-management-
profile default

WMM Traffic management profile "default"
----------------------------------------
Parameter              Value
---------              -----
Enable Shaping Policy  true
Voice Share            20 %
Video Share            50 %
Best-effort Share      20 %
Background Share       10 %
```

After you configure the WMM bandwidth management profile, apply it to the virtual AP profile.

```
(config) #wlan virtual-ap default
(Virtual AP profile "default") #wmm-traffic-management-profile default
```

9. Enable multicast shaping on the firewall.

```
(host) (config) #firewall shape-mcast
(host) (config) #show firewall

Global firewall policies
------------------------
Policy                                        Action    Rate      Slot/Port
------                                        ------    ----      ---------
Enforce TCP handshake before allowing data    Disabled
Prohibit RST replay attack                    Disabled
...
...
Multicast automatic shaping                   Enabled
Clear Sessions on Role Update                 Disabled
Session mirror IPSEC                          Disabled
```

Working with QoS for Voice and Video

QoS settings for voice and video applications are configured when you configure firewall roles and policies.

Understanding VoIP Call Admission Control Profile

VoIP call admission control prevents any single AP from becoming congested with voice calls. You configure call admission control options in the VoIP Call Admission Control profile, which you apply to an AP group or a specific AP.

You can use the WebUI or CLI to configure a VoIP Call Admission Control profile.

In the WebUI

1. Navigate to the **Configuration > AP Configuration** page. Select either **AP Group** or **AP Specific**.
   - If you select **AP Group**, click **Edit** for the AP group name for which you want to configure VoIP CAC.
   - If you select **AP Specific**, select the name of the AP for which you want to configure VoIP CAC.
2. In the **Profiles** list, expand the **QoS** menu, then select the **VoIP Call Admission Control** profile.
3. In the **Profile Details** window pane, click the **VoIP Call Admission Control profile** drop-down list and select the profile you want to edit.
   -or-
   To create a new profile, click the **VoIP Call Admission Control** profile drop-down list and select **New**. Enter a new profile name in the field to the right of the drop-down list. You cannot use spaces in VoIP profile names.
4. Configure your desired **VoIP Call Admission Control** profile settings. Table 190 describes the parameters you can configure in this profile:

Table 190: *VoIP Call Admission Control Configuration Parameters*

Parameter	Description
VoIP Call Admission Control	Select the **Voip Call Admission Control** checkbox to enable Wi-Fi VoIP Call Admission Control features.
VoIP Bandwidth based CAC	Select the **VoIP Bandwidth based CAC** checkbox to enable call admission controls based upon bandwidth. If this option is not selected, call admission controls are based on call counts.
VoIP Call Capacity	The maximum number of simultaneous calls that the AP radio can handle. You can use the bandwidth calculator in the WebUI to calculate the call capacity. To access the bandwidth calculator, navigate to **Configuration > Management > Bandwidth Calculator**. Default value: 10.
VoIP Bandwidth Capacity (kbps)	Enter a rate from 1 to 600000 (inclusive) to specify the maximum bandwidth rate that a radio can handle, in kbps. Default value is 2000 kbps.
VoIP Call Handoff Reservation	Specify the percentage of call capacity reserved for mobile VoIP clients on an active call. Default value is 20%.
VoIP Send SIP 100 Trying	The SIP invite call setup message is time-sensitive, as the originator retries the call as quickly as possible if it does not proceed. You can direct the controller to immediately reply to the call originator with a “SIP 100 - trying” message to indicate that the call is proceeding and to avoid a possible timeout. This is useful in conditions where the SIP invite may be redirected through a number of servers before reaching the controller.

Parameter	Description
	Select the **VoIP Send SIP 100 Trying** checkbox to send "SIP 100-trying" messages to a call originator to indicate that the call is proceeding. This is a useful option when the SIP invite is directed through many servers before reaching the controller.
VoIP Disconnect Extra Call	In the VoIP Call Admission Control (CAC) profile, you can limit the number of active voice calls allowed on a radio. This feature is disabled by default. When you enable the disconnect extra call feature, the system monitors the number of active voice calls, and if the defined threshold is reached, any new calls are disconnected. The AP denies association requests from a device that is on call. To enable this feature, select the **VoIP Disconnect Extra Call** checkbox. You also need to enable call admission control in this profile.
VOIP TSPEC Enforcement	A WMM client can send a Traffic Specification (TSPEC) signaling request to the AP before sending traffic of a specific AC type, such as voice. You can configure the controller so that the TSPEC signaling request from a client is ignored if the underlying voice call is not active. This feature is disabled by default. If you enable this feature, you can also configure the time duration within which the station should start the voice call after sending the TSPEC request (the default is one second). Select the **VoIP TSPEC Enforcement** checkbox to validate TSPEC requests for CAC.
VOIP TSPEC Enforcement Period	Select the maximum time, in seconds, for the station to start the call after the TSPEC request.
VoIP Drop SIP Invite and send status code (client)	Click the **VoIP Drop SIP Invite and send status code (client)** drop-down list and select one of the following status codes to be sent back to the client: • **480**: Temporary Unavailable • **486**: Busy Here • **503**: Service Unavailable • **none**: Don't send SIP status code
VoIP Drop SIP Invite and send status code (server)	Click the **VoIP Drop SIP Invite and send status code (client)**drop-down list and select one of the following status codes to be sent back to the server: • **480**: Temporary Unavailable • **486**: Busy Here • **503**: Ser vice Unavailable • **none**: Don't send SIP status code

5. Click **Apply**.

In the CLI

Use the following command:

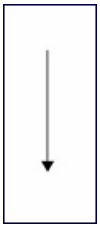
```
wlan voip-cac-profile <profile>
   bandwidth-cac
   bandwidth-capacity <bandwidth-capacity>
   call-admission-control
   call-capacity
   call-handoff-reservation <percent>
   disconnect-extra-call
   send-sip-100-trying
   send-sip-status-code client|server <code>
   wmm-tspec-enforcement
   wmm-tspec-enforcement-period <seconds>
```

Understanding Wi-Fi Multimedia

Wi-Fi Multimedia (WMM), is a Wi-Fi Alliance specification based on the IEEE 802.11e wireless Quality of Service (QoS) standard. WMM works with 802.11a, b, g, and n physical layer standards.

WMM supports four access categories (ACs): voice, video, best effort, and background. Table 191 shows the mapping of the WMM access categories to 802.1p priority values. The 802.1p priority value is contained in a two-byte QoS control field in the WMM data frame.

Table 191: *WMM Access Category to 802.1p Priority Mapping*

Priority	802.1p Priority	WMM Access Category
Lowest	1	Background
	2	
	0	Best effort
	3	
	4	Video
	5	
	6	Voice
Highest	7	

In non-WMM, or hybrid environments where some clients are not WMM-capable, Dell uses voice and best effort to prioritize traffic from these clients.

Unscheduled Automatic Power Save Delivery (U-APSD) is a component of the IEEE 802.11e standard that extends the battery life on voice over WLAN devices. When enabled, clients trigger the delivery of buffered data from the AP by sending a data frame.

For the environments in which the wireless clients support WMM, you can enable both WMM and U-APSD in the SSID profile.

Enabling WMM

You can use the WebUI or CLI to enable WMM for wireless clients.

In the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. Select either the AP Group or AP Specific tab. Click **Edit** for the AP group or AP name.
3. In the **Profiles** list, select **Wireless LAN**, then **Virtual AP**, then the applicable virtual AP profile. Select the **SSID** profile.
4. In the **Profile Details**, select the **Advanced** tab.
5. Select the **Wireless Multimedia (WMM)** option. Or, select the **Wireless Multimedia U-APSD (WMM-UAPSD) Powersave** option if you want to enable WMM in power save mode.
6. Click **Apply**.

In the CLI

Use the following command:

```
wlan ssid-profile <profile> wmm
wlan ssid-profile <profile> wmm-uapsd
```

Configuring WMM AC Mapping

The IEEE 802.11e standard defines the mapping between WMM ACs and Differentiated Services Codepoint (DSCP) tags. The WMM AC mapping commands allow you to customize the mapping between WMM ACs and DSCP tags to prioritize various traffic types. You apply and configure WMM AC mappings to a WMM-enabled SSID profile.



Ensure that you enable WMM for legacy APs for the mapping to take effect. For 802.11n APs, ensure that you enable either WMM or high throughput.

DSCP classifies packets based on network policies and rules, not priority. The configured DSCP value defines per hop behaviors (PHBs). The PHB is a 6-bit value added to the 8-bit Differentiated Services (DS) field of the IP packet header. The PHB defines the policy and service applied to a packet when traversing the network. You configure these services in accordance with your network policies. Table 192 shows the default WMM AC to DSCP decimal mappings and the recommended WMM AC to DSCP mappings.

Table 192: *WMM Access Category to DSCP Mappings*

DSCP Decimal Value	WMM Access Category
8	Background
16	
0	Best effort
24	
32	Video
40	
48	Voice
56	

By customizing WMM AC mappings, both the controller and AP maintain a customized WMM AC mapping table for each configured SSID profile. All packets received are matched against the entries in the mapping table and prioritized accordingly. The mapping table contains information for upstream (client to AP) and downstream (AP to client) traffic.



In earlier releases, the default mappings exist for all SSIDs. After you customize a WMM AC mapping and apply it to the SSID, the controller overwrites the default mapping values and uses the configured values . If a controller is upgraded to 6.2 from an older version, the default and the user configured WMM-DSCP mappings in the existing SSID profiles are retained. There are no default mappings for a newly created SSID profile and for a factory default controller running 6.2 image.

When planning your mappings, make sure that any immediate switch or router does not have conflicting 802.1p or DSCP configurations/mappings. If this occurs, your traffic may not be prioritized correctly.

To view the mapping settings, use the following command:

```
show wlan ssid-profile <profile>
```

Using the WebUI to map between WMM AC and DSCP

1. Navigate to the **Configuration > Wireless > AP Configuration** page.
2. Select either the **AP Group** or **AP Specific** tab. Click **Edit** for the AP group or AP name.

3. In the **Profiles** list, select **Wireless LAN**, then **Virtual AP**, then the applicable virtual AP profile. Select the SSID profile.
4. In the **Profile Details**, select the **Advanced** tab.
5. Scroll down to the **Wireless Multimedia (WMM)** option. Select this option.
6. Modify the DSCP mapping settings, as needed:
   - DSCP mapping for WMM voice AC—DSCP used to map voice traffic
   - DSCP mapping for WMM video AC—DSCP used to map video traffic
   - DSCP mapping for WMM best-effort AC—DSCP used to map best-effort traffic
   - DSCP mapping for WMM background AC—DSCP used to map background traffic
7. Click **Apply**.

The following enhancements have been made to the WMM-DSCP mapping functionality:

- When a controller is upgraded to the 6.2 version from an older version, the default and the user configured WMM-DSCP mappings in the existing SSID profiles are retained.
- Default mappings are not there for a newly created SSID profile and for a factory default controller a running 6.2 image.
- If the mapping has no value, the original DSCP for upstream traffic is retained.
- The maximum number of values that can be configured for WMM-DSCP is 8.
- For the upstream traffic, if the mapping exists and incoming DSCP value matches one of the mapped values, then the DSCP value is retained.
- For the upstream traffic, if the mapping exists and incoming DSCP value does not match any of the mapped values, then the DSCP value is overwritten with the first value in the WMM- DSCP list
- For Wireless to Wireless Traffic: If the AC of the incoming packet has no mapping and the incoming DSCP value is mapped to a different AC, then the DSCP value is retained and WMM priority is changed to the corresponding AC where incoming DSCP is mapped.

Using the CLI to map between WMM AC and DSCP

Use the following commands:

```
wlan ssid-profile <profile>
   wmm-be-dscp <best-effort>
   wmm-bk-dscp <background>
   wmm-vi-dscp <video>
   wmm-vo-dscp <voice>
```

Configuring DSCP Priorities

You can configure DSCP priorities for WMM packets in the following ways:

- configure the DSCP mappings in the SSID profile
- set a ToS value in the ACL
- set the ToS value and the 802.1p priority in the ACL

Setting a ToS value in the ACL overrides the default DSCP mappings configured in the SSID profile. Configuring a DSCP priority in both the L2 and L3 header prioritizes the WMM packets with the higher value.

For example, you can have different ToS values set for different voice traffic in a network. To prioritize all of them in the voice queue, we can set the 802.1p priority to voice.

Consider a deployment where Cisco Softphone, Lync, and Scopia are configured with the following DSCP :

- Cisco Softphone - DSCP 46
- Lync - DSCP 44

- Scopia - DSCP 42

In the absence of doing anything, all of the DSCP above would map into the Video queue. To map all the traffic into voice queue, you can use the following ACL configuration:

```
wlan ssid-profile VOICE
   wmm-vo-dscp 46
ip access-list session VOICE
   any destination [LYNC_SERVER] [LYNC_PORTS] permit tos 44 dot1p-priority 6
   any destination [SCOPiA_SERVER] [SCOPIA _PORTS] permit tos 42 dot1p-priority 6
```



You must know the ports on which each traffic is sent so that the correct traffic is identified.

Configuring Dynamic WMM Queue Management

Traditional wireless networks provide all clients with equal bandwidth access. However, delays or reductions in throughput can adversely affect voice and video applications, resulting in disrupted VoIP conversations or dropped frames in a streamed video. Thus, data streams that require strict latency and throughput need to be assigned higher traffic priority than other traffic types.

The Wi-Fi Alliance defined the Wi-Fi Multimedia (WMM) standard in response to industry requirements for Quality of Service (QoS) support for multimedia applications for wireless networks. This is defined as per the IEEE 802.11e standards.

WMM requires:

- the access point be Wi-Fi Certified and has WMM enabled
- the client device be Wi-Fi Certified
- the application support WMM

Enhanced Distributed Channel Access

WMM provides media access prioritization through Enhanced Distributed Channel Access (EDCA). EDCA defines four access categories (ACs) to prioritize traffic: voice, video, best effort, and background. These ACs correspond to 802.1p priority tags, as shown in Table 193.

Table 193: *WMM Access Categories and 802.1p Tags*

WMM Access Category	Description	802.1p Tag
Voice	Highest priority	7, 6
Video	Prioritize video traffic above other data traffic	5, 4
Best Effort	Traffic from legacy devices or traffic from applications or devices that do not support QoS	0, 3
Background	Low priority traffic (file downloads, print jobs)	2, 1

While the WMM ACs designate specific types of traffic, you can determine the priority of the ACs. For example, you can choose to give video traffic the highest priority. With WMM, applications assign data packets to an AC. In the client, the data packets are then added to one of the transmit queues for voice, video, best effort, or background.

WMM is an extension to the Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) protocol's Distributed Coordination Function (DCF). The collision resolution algorithm responsible for traffic prioritization depends on the following configurable parameters for each AC:

- arbitrary inter-frame space number (AIFSN)

- minimum and maximum contention window (CW) size

For each AC, the backoff time is the sum of the AIFSN and a random value between 0 and the CW value. The AC with the lowest backoff time is granted the opportunity to transmit (TXOP). Frames with the highest-priority AC are more likely to get TXOP, because they tend to have the lowest backoff times (a result of having smaller AIFSN and CW parameter values). The value of the CW varies through time as the CW doubles after each collision up to the maximum CW. The CW is reset to the minimum value after successful transmission. In addition, you can configure the TXOP duration for each AC.

On the controller, you configure the AC priorities in the WLAN EDCA parameters profile. There are two sets of EDCA profiles you can configure:

- AP parameters affecting traffic from the AP to the client.
- STA parameters affecting traffic from the client to the AP.

Using the WebUI to configure EDCA parameters

Use the following procedure to define an Enhanced Distributed Channel Access (EDCA) profile for APs or for clients (stations).

1. Navigate to the **Configuration > AP Configuration** page. Select either the **AP Group** tab or **AP Specific** tab:
   - If you selected **AP Group**, click **Edit** for the AP group name for which you want to configure EDCA parameters.
   - If you selected **AP Specific**, select the name of the AP for which you want to configure EDCA parameters.
2. Under **Profiles**, select the **Wireless LAN**, then **Virtual AP**. In the Virtual AP list, select the appropriate virtual AP.
3. Expand the **SSID** profile. Select the **EDCA Parameters Station** or **EDCA Parameters AP** profile.
4. Configure your desired EDCA Profile Parameters. Table 194 describes the parameters you can configure in this profile.

Table 194: *EDCA Parameters Station and EDCA Parameters AP Profile Settings*

Parameter	Description
Best Effort	Set the following parameters to define the best effort queue: • **aifsn**: arbitrary inter-frame space number. Range: 1-15. • **ecw-max**: the exponential (n) value of the maximum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 1-15. • **ecw-min**: the exponential (n) value of the minimum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 0-15. • **txop**: transmission opportunity, in units of 32 microseconds. Divide the desired transmission duration by 32 to determine the value to configure. For example, for a transmission duration of 3008 microseconds, enter 94 (3008/32). Range: 0-2047. • **acm**: this parameter specifies mandatory admission control. With a value of **1**, the client reserves the access category through traffic specification (TSPEC) signaling. A value of **0** disables this option.
Background	Set the following parameters to define the background queue: • **aifsn**: arbitrary inter-frame space number. Range: 1-15. • **ecw-max**: the exponential (n) value of the maximum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 1-15. • **ecw-min**: the exponential (n) value of the minimum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 0-15. • **txop**: transmission opportunity, in units of 32 microseconds. Divide the desired transmission duration by 32 to determine the value to configure. For example, for a transmission duration of 3008 microseconds, enter 94 (3008/32). Range: 0-2047. • **acm**: this parameter specifies mandatory admission control. With a value of **1**, the client reserves the access category through traffic specification (TSPEC) signaling. A value of **0** disables this option.
Video	Set the following parameters to define the background queue: • **aifsn**: trbitrary inter-frame space number. Range: 1-15. • **ecw-max**: The exponential (n) value of the maximum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 1-15. • **ecw-min**: the exponential (n) value of the minimum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 0-15. • **txop**: transmission opportunity, in units of 32 microseconds. Divide the desired transmission duration by 32 to determine the value to configure. For example, for a transmission duration of 3008 microseconds, enter 94 (3008/32). Range: 0-2047. • **acm**: this parameter specifies mandatory admission control. With a value of **1**, the client reserves the access category through traffic specification (TSPEC) signaling. A value of **0** disables this option.
Voice	Set the following parameters to define the background queue: • **aifsn**: trbitrary inter-frame space number. Range: 1-15. • **ecw-max**: the exponential (n) value of the maximum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 1-15. • **ecw-min**: the exponential (n) value of the minimum contention window size, as expressed by 2^n-1. A value of 4 computes to 2^4-1 = 15. Range: 0-15. • **txop**: transmission opportunity, in units of 32 microseconds. Divide the desired transmission duration by 32 to determine the value to configure. For example, for a transmission duration of 3008 microseconds, enter 94 (3008/32). Range: 0-2047. • **acm**: this parameter specifies mandatory admission control. With a value of **1**, the client reserves the access category through traffic specification (TSPEC) signaling. A value of **0** disables this option.

5. Click **Apply**.

Using the CLI to configure EDCA parameters

Use the following commands:

```
wlan edca-parameters-profile {ap|station} <profile>
{background | best-effort | video | voice}
```

```
[acm][aifsn <number>] [ecw-max <exponent> [ecw-min <exponent>] [txop <number>]
```

To associate the EDCA profile instance to a SSID profile:

```
wlan ssid-profile <profile>
edca-parameters-profile {ap|sta} <profile>
```

Enabling WMM Queue Content Enforcement

WMM queue content enforcement is a firewall setting that you can enable to ensure that the voice priority is used for voice traffic. When you enable this feature, if traffic to or from the user is inconsistent with the associated QoS policy for voice, the traffic is reclassified to best effort and data path counters incremented. If TSPEC admission were used to reserve bandwidth, then TSPEC signaling informs the client that the reservation is terminated.

You can use the WebUI or CLI to enable WMM queue content enforcement.

In the WebUI

1. Navigate to the **Configuration > Advanced Services > Stateful Firewall** page.
2. Select **Enforce WMM Voice Priority Matches Flow Content**.
3. Click **Apply**.

In the CLI

Use the following command:

```
firewall wmm-voip-content-enforcement
```

Unified Communication and Collaboration

This section describes the Unified Communication and Collaboration (UCC) feature. The Unified Communications Manager (UCM) is the core solution component of this feature. UCC addresses the onslaught of mobile devices that use voice, video, and collaboration applications. UCC solution reduces the cost of infrastructure for enterprise communication and collaboration.



The UCC feature requires the PEFNG license.

UCC continues to support all existing functionality provided by ArubaOS 6.3.x. This section includes the following sub-sections:

- Microsoft® Lync Visibility and Granular QoS Prioritization on page 899
- UCC Dashboard in the WebUI on page 908
- Viewing UCC Information on page 912
- UCC Troubleshooting on page 912
- UCC-AirWave Integration on page 912
- UCC Call Quality Metrics on page 913
- Changes to Call Admission Control on page 914
- UCC Limitations on page 915

Microsoft® Lync Visibility and Granular QoS Prioritization

This release of ArubaOS provides a seamless user experience for Microsoft® Lync users using voice or video calls, desktop sharing, and file transfer in a wireless environment. Microsoft Lync is an enterprise solution for UCC. It provides support for voice, video, desktop-sharing, and file-transfer.

NOTE

Microsoft Lync uses SIP over TLS for call signaling.

ArubaOS provides value added services such as prioritization of Lync sessions, call quality metrics, and visibility by implementing Lync Application Layer Gateway (ALG). This solution also provides a dedicated visibility and troubleshooting framework that allows network administrators to fine-tune and troubleshoot Lync traffic flow in the network.

As Microsoft Lync deployments are more widely implemented on wireless networks, it is important to provide Quality of Service (QoS) for Lync voice or video calls, desktop sharing, and file transfer so that there is no visible difference in the user experience between wireless and wired networks. Lync ALG offers an enriched solution in terms of QoS, scalable voice, video, desktop-sharing, and file-transfer. The ALG based solution provides the following value-added services:

- **Call Quality Metrics**: Call quality details such as Mean Opinion Score (MOS), delay, jitter, and packet loss.
- **Call Priority**: Call priority is provided for all Lync sessions irrespective of Call Admission Control (CAC) limit. In ArubaOS 6.4, Lync sessions are prioritized based on session-specific requirements. Voice calls get highest priority, followed by video and desktop sharing. File-transfer gets the least priority.
- **Call and Client information**: Details about Lync call types and statistics through CLI. The commands are discussed later in this section.
- **Deterministic Solution**: A dedicated visibility and troubleshooting framework that allows network administrators to fine-tune and troubleshoot Lync traffic flow in the network. This solution provides an enriched performance which uses Deep Packet Inspection (DPI).
- **Call Admission Control**: In ArubaOS 6.4, Lync sessions do not come under the purview of call count based CAC and bandwidth based CAC. For more information, see Changes to Call Admission Control on page 914.

To take advantage of ArubaOS 6.4 UCC Lync ALG, it is recommended to use the Microsoft Lync SDN API. ArubaOS 6.4 supports all versions of Lync SDN API up to version 2.0. Microsoft Lync SDN API works with Microsoft Lync server to export details about voice or video calls, desktop-sharing, and file-transfer to Dell controller's Web server. The communication between the Lync SDN API and Web server occurs over HTTP or HTTPS.

This section includes the following sub-sections:

- Lync ALG Compatibility Matrix on page 900
- Configuration Prerequisites on page 901
- Configuring Lync ALG on page 901
- Viewing Lync ALG Statistics using the CLI
- Viewing Lync ALG Statistics Using the WebUI on page 906
- Troubleshooting Lync ALG Issues on page 907

Lync ALG Compatibility Matrix

The following table lists the Lync clients that support voice, video, desktop sharing, and file transfer applications in ArubaOS 6.4:

Table 195: *Compatibility Matrix*

Lync Client	Lync Server 2010	Lync Server 2013
Android	No	Yes
iOS	No	Yes
OS X (Mac)	Yes	Yes
Windows	Yes	Yes

Configuration Prerequisites

- Microsoft Lync server supporting Lync SDN API versions up to 2.0.
- Dell controller running ArubaOS 6.4.



If your setup does not have a Lync SDN API, use **Media Classification** as described in Understanding Extended Voice and Video Features on page 915.

Configuring Lync ALG

This section describes the procedures to configure Lync ALG on the controller:

- Configuring Lync Listening Port on page 901
- Configuring Lync ALG Status on page 902
- Dynamically Open Firewall for UCC Clients using STUN on page 902
- Configuring Per User Role Lync Call Prioritization on page 903
- Disable Media Classification on page 905

When upgrading from ArubaOS 6.x to 6.4:

- Lync ALG is enabled by default.
- If media classification is configured before upgrading to ArubaOS 6.4, disable media classification.

Configuring Lync Listening Port

Configure the port number on which Microsoft Lync SDN API sends HTTP or HTTPS call information (XML) messages to Dell controller.



Before you configure Lync listening port, disable **classify-media**. To disable **classify-media**, see Disable Media Classification on page 905.

Using the WebUI

1. Navigate to the **Configuration > Management > General** page.
2. Under the **Configure Lync** section, select the **HTTP** or **HTTPS** protocol from the drop-down list and enter the port number in the **Web lync listening port** text box.

   The port range is from 1024 to 65535.



The **Web lync listening port** is automatically permitted by the firewall. The user does not have to explicitly define a firewall policy to permit this port.

3. Click **Apply**.

Using the CLI

Use the following command:

```
(host) (config) #web-server
```

Listen Lync XML messages on HTTP:

```
(host) (Web Server Configuration) #web-lync-listen-port http <listen-port>
```

Or

Listen Lync XML messages on HTTPS:



Before configuring the controller to receive Lync SDN API messages using HTTPS, a server certificate must be generated and installed on the controller. Server certificate can be generated either by the controller or Certificate Authority (CA). For more information, see Obtaining a Server Certificate on page 796.

```
(host) (Web Server Configuration) #web-lync-listen-port https <listen-port>
```

To verify if the port is automatically permitted by the firewall, use the following command:

```
(host) #show firewall-cp
```

Configuring Lync ALG Status

Configure the controller to read Secure SIP signaling messages sent by the Lync clients on port 5061. You can enable or disable Stateful SIPS processing using the following CLI commands. This is enabled by default.



Before you configure Lync ALG status, disable **classify-media**. To disable **classify-media**, see Disable Media Classification on page 905.

Enabling Lync ALG

```
(host) (config) #no firewall disable-stateful-sips-processing
```

Disabling Lync ALG

```
(host) (config) #firewall disable-stateful-sips-processing
```

Dynamically Open Firewall for UCC Clients using STUN

Prior to ArubaOS 6.4, the administrator explicitly added ACLs in the user role to allow Lync traffic on the controller. Starting from ArubaOS 6.4, the controller automatically allows firewall sessions for Lync voice and video calls. Firewall sessions for Lync desktop-sharing and file-transfer are not allowed. The administrator should manually open a range of TCP ports under the user role to allow Lync desktop-sharing and file-transfer traffic.To allow a specific range of ports in the user role, refer the Microsoft Technet article which describes the port ranges used by Lync clients and servers.

Before media transmission, a Lync client initiates a Session Traversal Utilities for NAT (STUN) connectivity check. Sessions created by STUN are subjected to media classification that classifies the media as Real-time Transport Protocol (RTP) or non-RTP. The firewall automatically allows the RTP session on the controller and denies the non-RTP sessions. For the controller to accept STUN messages, you must allow ICE-STUN based firewall traversal on the controller and allow UDP 3478 and TCP 443 ports in the user role.

Allowing ICE-STUN

To allow ICE-STUN based firewall traversal, issue the following CLI command:

```
(host) (config) #firewall allow-stun
```

Allowing UDP Port 3478

STUN uses UDP port 3478. To allow UDP port 3478 in the user role, issue the following CLI commands.

```
(host) (config) #user-role <STRING>
(host) (config-role) #ip access-list session stun
(host) (config-sess-stun)#any any udp 3478 permit
```

Allowing TCP Port 443

HTTP Secure (HTTPS) uses TCP port 443. To allow TCP port 443 in the user role, issue the following CLI commands.

```
(host) (config) #user-role <STRING>
(host) (config-role) #ip access-list session https-acl
(host) (config-sess-stun)#any any svc-https permit
```

Configuring Per User Role Lync Call Prioritization

In ArubaOS 6.3.x, you can configure the UCC call prioritization system-wide only. For example, Lync voice, video, and collaboration applications can be configured system-wide on the controller. Starting from ArubaOS 6.4, an administrator can configure Lync call prioritization on a per user-role basis. With this feature, you can have one set of users have priority on real-time media traffic over another set of users. An administrator can configure the per user-role Lync call prioritization based on the deployment needs.

Important Points to Remember

If two clients in an active call are in different user-roles and traffic prioritization, the traffic prioritization is based on the following order:

1. Voice
2. Video
3. Best-effort
4. Background

The above is applicable for both; the caller and called parties.

Example

Client 1 (C1) is assigned user-role 1 with voice priority enabled. Client 2 (C2) is assigned user-role 2 with voice priority disabled. When C1 makes a voice call to C2, both parties have prioritized voice call.



Traffic prioritization may not apply if it is a conference call or the caller and called parties are in multiple controllers.

You can configure the per user-role UCC call prioritization for Lync ALG using the WebUI or CLI.

Using the WebUI

Configure the Lync traffic control profile:

1. Navigate to the **Configuration > Advanced Services > All Profiles** page.
2. Select **Other Profiles** to expand the **Other Profiles** section.
3. Click the **Traffic Control Prioritization** profile.
4. Under the **Traffic Control Prioritization Profile** section, enter the profile name and click **Add**.
5. Click the newly created profile.
6. Select the appropriate check box to prioritize Lync traffic. See .



By default, Lync ALG prioritizes all the four application types.

Table 196: *Lync ALG Traffic Priority Parameters*

Traffic Control Parameter	Description
Prioritize voice	Prioritizes voice sessions by Lync ALG.
Prioritize video	Prioritizes video sessions by Lync ALG.
Prioritize desktop-sharing	Prioritizes desktop sharing sessions by Lync ALG.
Prioritize file-transfer	Prioritizes file transfer sessions by Lync ALG.

Link the newly created Lync traffic control profile to the user-role:

1. Navigate to the **Configuration > Security > Access Control > User Roles** page.
2. Select an existing user role, and click **Edit**.
3. Under the **Misc. Configuration** section, select the newly created Lync traffic control profile from the **Traffic Control Profile** drop-down list.

Using the CLI

Configure the Lync traffic control profile:

```
(host) (config) #app lync traffic-control <profile-name>
(host) (Traffic Control Prioritization Profile "default") #prioritize-voice
(host) (Traffic Control Prioritization Profile "default") #prioritize-video
(host) (Traffic Control Prioritization Profile "default") #prioritize-desktop-sharing
(host) (Traffic Control Prioritization Profile "default") #prioritize-file-transfer
```

To verify the configuration, use the following command:

```
(host) #show ucc configuration traffic-control lync <profile-name>
```

Link the newly created Lync traffic control profile to the user-role.

```
(host) (config) #user-role <STRING>
(host) (config-role) #traffic-control-profile <STRING>
```

Recommended DSCP Mapping for Lync Traffic in Dell Controller

The following DSCP values for Lync ALG are recommended:

Lync Application	DSCP Mapping
Voice	56
Video and desktop sharing	40
File transfer	24 (Best-effort)

You can configure the DSCP mappings in the SSID profile using the following CLI command:

```
(host)(config) #wlan ssid-profile Lync_ALG
(host) (SSID Profile "Lync_ALG") #wmm
(host) (SSID Profile "Lync_ALG") #wmm-vo-dscp 56
(host) (SSID Profile "Lync_ALG") #wmm-vi-dscp 40
(host) (SSID Profile "Lync_ALG") #wmm-be-dscp 24
```

Disable Media Classification



Media classification is not supported when clients are accessed through a network address translation (NAT).

Media classification should not be configured on session ACL for Secure SIP used by Lync clients. The following example verifies if media classification is configured on session ACL that is associated with the user-role, "employee":

```
(host) #show rights employee

Derived Role = 'employee'
Up BW:No Limit   Down BW:No Limit
L2TP Pool = default-l2tp-pool
PPTP Pool = default-pptp-pool
Periodic reauthentication: Disabled
ACL Number = 64/0
Max Sessions = 65535

access-list List
----------------
Position  Name       Type     Location
--------  ----       ----     --------
1         employee   session

employee
---------
Priority  Source  Destination  Service       Action  TimeRange  Log
--------  ------  -----------  -------       ------  ---------  ---
1         any     any          svc-sips      permit

Expired  Queue  TOS  8021P  Blacklist  Mirror  DisScan  ClassifyMedia
-------  -----  ---  -----  ---------  ------  -------  -------------
         High                                         Yes

IPv4/6
------
4

Expired Policies (due to time constraints) = 0
```

Under **ClassifyMedia** column, **Yes** indicates media classification is configured. To disable it, you must first delete the ACL. Use the following commands:

```
(host) (config) #ip access-list session employee
(host) (config-sess-employee) #no any any svc-sips permit
```

You must add the rule **any any svc-sips permit** back to the ACL without the **classify-media** parameter:

```
(host) (config-sess-employee) #any any svc-sips permit
```

Viewing Lync ALG Statistics using the CLI

This section describes the procedures to view Lync ALG statistics using the CLI.



For detailed command parameters, see the *ArubaOS 6.4 CLI Reference Guide*.

- Viewing the list of Lync Clients on page 906
- Viewing Call Detail Record for Lync Calls on page 906
- Viewing Call Quality for Lync Calls on page 906

- Viewing Lync Call Trace Buffer on page 906

Viewing the list of Lync Clients

Use the following command to display details of clients that are actively using Lync. An entry is created for clients that have actively participated in voice, video, desktop-sharing, or file-sharing sessions.

```
(host) #show ucc client-info app lync
```

Viewing Call Detail Record for Lync Calls

Use the following command to view the Call Detail Record for Lync calls on the controller. This command displays the last 128 call records for W-600 Series controller platform and 512 call records for rest of the controller platforms.

```
(host) #show ucc call-info cdrs app lync
```

Viewing Call Quality for Lync Calls

Use the following command to view the call quality information for Lync voice and video calls.

```
(host) #show ucc call-info cdrs detail
```

Viewing Lync Call Trace Buffer

Use the following command to display the Lync message trace buffer for the first 256 events. Events such as establishing voice, video, desktop sharing, and file transfer are recorded.

```
(host) #show ucc trace-buffer lync
```

Viewing Lync ALG Statistics Using the WebUI

This section describes the procedures to view Lync ALG statistics using the WebUI.

- Viewing Voice Status on page 906
- Viewing Call Performance Report on page 906
- Viewing Call Density Report on page 906
- Viewing Call Detail Report on page 907
- Viewing Voice Client Call Statistics on page 907
- Viewing Voice Client HandOff Information on page 907
- Viewing Voice Client Troubleshooting Information on page 907

Viewing Voice Status

To view the status of Lync calls:

1. Navigate to the **Monitoring > VOICE > Voice Status** page.
2. Select **Lync** from the **Protocol** drop-down list.

Viewing Call Performance Report

This report displays the performance of voice calls of Lync clients connected to the controller. You can filter the report based on AP IP address, BSSID, Client Extension, ESSID, or the VOIP protocol type. To view the performance of Lync calls:

1. Navigate to the **Monitoring > VOICE > Call Performance Report** page.
2. Select **Lync** from the **Protocol** drop-down list.

Viewing Call Density Report

To view the call density report of Lync calls:

1. Navigate to the **Monitoring > VOICE > Call Density Report** page.

2. Select **Lync** from the **Protocol** drop-down list.

Viewing Call Detail Report

This report displays detailed call records of Lync clients. To view the report:

1. Navigate to the **Monitoring > VOICE > Call Detail Report** page.
2. Select **Lync** from the **Protocol** drop-down list.

Viewing Voice Client Call Statistics

To view call statistics of a Lync client:

1. Navigate to the **Monitoring > VOICE > Voice Clients** page.
2. Select **Lync** from the **Protocol** drop-down list.
3. Select a client from the list of client IP and click **View Call Statistics** button.

Viewing Voice Client HandOff Information

To view the handoff information of a Lync client:

1. Navigate to the **Monitoring > VOICE > Voice Clients** page.
2. Select **Lync** from the **Protocol** drop-down list.
3. Select a client from the list of client IP and click **HandOff Information**.

Viewing Voice Client Troubleshooting Information

To view troubleshooting information of a Lync client:

1. Navigate to the **Monitoring > VOICE > Voice Clients** page.
2. Select **Lync** from the **Protocol** drop-down list.
3. Select a client from the list of client IP and click **Troubleshooting**.

Troubleshooting Lync ALG Issues

The following sections describe the CLI commands to troubleshoot Lync ALG issues.

Enabling Lync ALG Debug Logs

Lync ALG related debug logs are available under logs. Use the following command to enable this:

```
(host) (config) #logging level debugging user process stm subcat voice
```

Viewing Lync ALG Debug Logs

To view the Lync ALG debug logs, use the following command:

```
(host) (config) #show log user all
Jul 18 15:33:09 :503188: <DBUG> |stm| |voice| vm_lync_create_call: mac(1c:ab:a7:2d:75:6b) num_
sessions(0) curr_session(0x1064f144)

Jul 18 15:33:09 :503188: <DBUG> |stm| |voice| VM: vm_lync_create_call:1001 LYNC INFO: Headers
are 2b00b11f-71e3-40a5-a1bf-386dc9d49eb6 sip:user@lyncqa.com sip:user1@lyncqa.com

Jul 18 15:33:09 :503188: <DBUG> |stm| |voice| vm_lync_update_session_sdp:1869 -- vc
(1c:ab:a7:2d:75:6b)

Jul 18 15:33:09 :503188: <DBUG> |stm| |voice| VM: vm_lync_update_session_sdp:1961 LYNC INFO:
copied 1 staus to call_info

Jul 18 15:33:09 :503162: <DBUG> |stm| |voice| VM: vm_lync_update_session_sdp 1963: Tx params
changed
```

```
Jul 18 15:33:09 :503188: <DBUG> |stm| |voice| vm_lync_update_session_sdp:1869 -- vc
(1c:ab:a7:2d:75:6b)

Jul 18 15:33:09 :503188: <DBUG> |stm| |voice| VM: vm_lync_update_session_sdp:1961 LYNC INFO:
copied 1 staus to call_info

Jul 18 15:33:09 :503162: <DBUG> |stm| |voice| VM: vm_lync_update_session_sdp 1963: Tx params
changed
Jul 18 15:33:09 :503126: <DBUG> |stm| |voice| VM: vm_lync_create_call 1023: Session created
and inserted successfully for call id 2b00b11f-71e3-40a5-a1bf-386dc9d49eb6, 10.XX.XX.208

Jul 18 15:33:09 :503188: <DBUG> |stm| |voice| VM: vm_lync_idle_startdialog_req:301 LYNC INFO:
vm_lync_create_call is success..

Jul 18 15:33:09 :503188: <DBUG> |stm| |voice| VM: vm_lync_idle_startdialog_req:309 LYNC INFO:
vm_lync_create_call() is success..
```

UCC Dashboard in the WebUI

The UCC dashboard gives a complete view of the UCC deployment in the controller. The UCC dashboard has two levels of displaying statistics:

- UCC Dashboard Aggregated Display
- UCC Dashboard Per Client Display

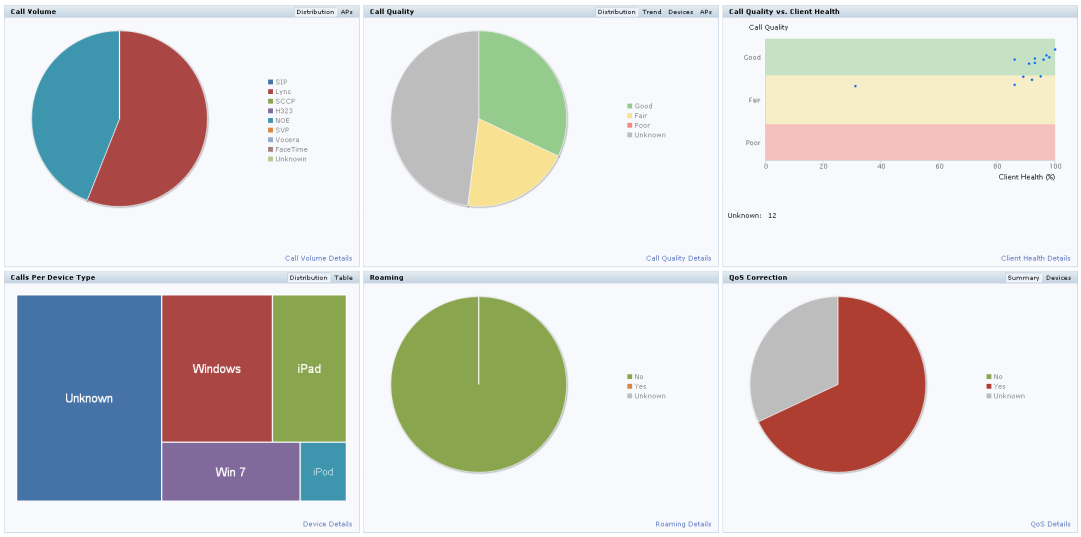
UCC Dashboard Aggregated Display

The UCC Dashboard Aggregated Display shows an aggregated view of the UCC calls made in the controller. The administrator can see a top level view of the call quality assessment, and further drill-down into a specific view based on the analysis required.

Chart View

Navigate to the **Dashboard > UCC** page. The **UCC** page displays the overall health (in graphical format) of the UCC deployment in the controller as shown in Figure 197.

Figure 197 *UCC Dashboard*



Each graphical section of the UCC dashboard is explained below:

- **Call Volume** – This graph displays the total number of calls made based on the UCC application type. For example, SIP, Lync, SCCP, H.323, NOE, SVP, VOCERA, and FaceTime.
- **Call Quality** – This graph displays the number of UCC calls categorized by the following call quality:

- **Good**
- **Fair**
- **Poor**
- **Unknown**: Voice or desktop-sharing sessions that do not originate or terminate on the home controller is categorized under **Unknown** call quality.

- **Call Quality vs. Client Health** - This graph displays the co-relation between the VoIP call quality and the VoIP client health of every UCC call.
- **Calls Per Device Type** – This graph displays the calls made per device type. For example, Windows 7, Mac OS X, iPhone, or Android.
- **Roaming** – Roaming status of UCC clients. The status can be:
  - **No** – Number of calls where the client did not roam to a new AP.
  - **Yes** - Number of calls where the client has roamed to a new AP.
  - **Unknown** - For external clients which are not registered with the controller is categorized as **Unknown**.
- **QoS Correction** – If the DSCP value of the Real-time Transport Protocol (RTP) packets sent by the client differs from the WMM-DSCP configured in the corresponding SSID profile definition, the controller corrects this value as per the SSID profile definition and classifies the call as QoS corrected. This graph displays the number of UCC calls where the controller has corrected the WMM-DSCP value for such calls. The QoS correction is categorized as:
  - **No** – No WMM-DSCP value correction.
  - **Yes** – WMM-DSCP value corrected by the controller.
  - **Unknown** – For external clients which are not registered with the controller is categorized as **Unknown**.

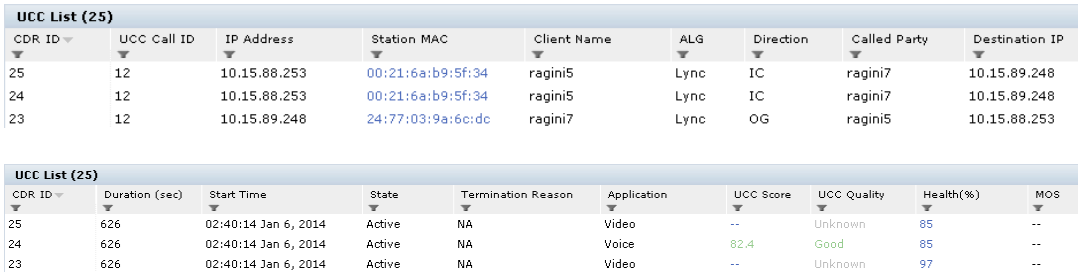
Details View

To display an aggregated list of all the UCC call data metrics in the controller, navigate to the **Dashboard > UCC** page of the WebUI and click any of the following hyperlinks:

- Call Volume Details
- Call Quality Details
- Client Health Details
- Device Details
- Roaming Details
- QoS Details

Figure 198 displays an aggregated list of all the UCC call data metrics in the controller.

Figure 198 *UCC List*

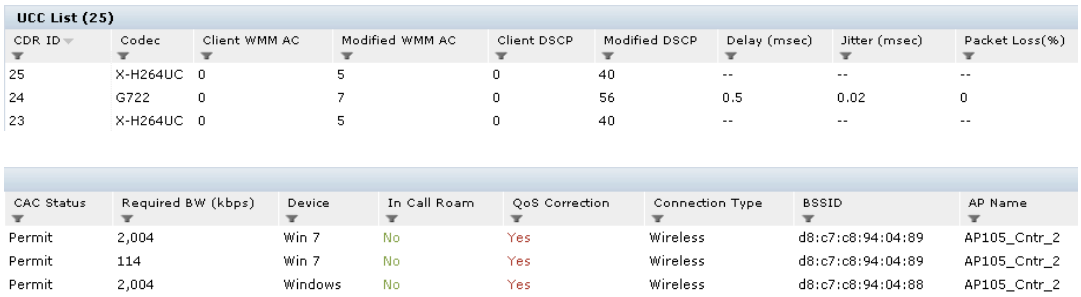


UCC List (25)

CDR ID	UCC Call ID	IP Address	Station MAC	Client Name	ALG	Direction	Called Party	Destination IP
25	12	10.15.88.253	00:21:6a:b9:5f:34	ragini5	Lync	IC	ragini7	10.15.89.248
24	12	10.15.88.253	00:21:6a:b9:5f:34	ragini5	Lync	IC	ragini7	10.15.89.248
23	12	10.15.89.248	24:77:03:9a:6c:dc	ragini7	Lync	OG	ragini5	10.15.88.253

UCC List (25)

CDR ID	Duration (sec)	Start Time	State	Termination Reason	Application	UCC Score	UCC Quality	Health(%)	MOS
25	626	02:40:14 Jan 6, 2014	Active	NA	Video	--	Unknown	85	--
24	626	02:40:14 Jan 6, 2014	Active	NA	Voice	82.4	Good	85	--
23	626	02:40:14 Jan 6, 2014	Active	NA	Video	--	Unknown	97	--



UCC Dashboard Per Client Display

Under the **Dashboard > Clients** page of the WebUI, clicking the client IP hyperlink displays the details page of the client. Click the **UCC** tab. This tab displays an aggregated list of UCC call data metrics of a client. See Figure 199.

Figure 199 *UCC Client Page*

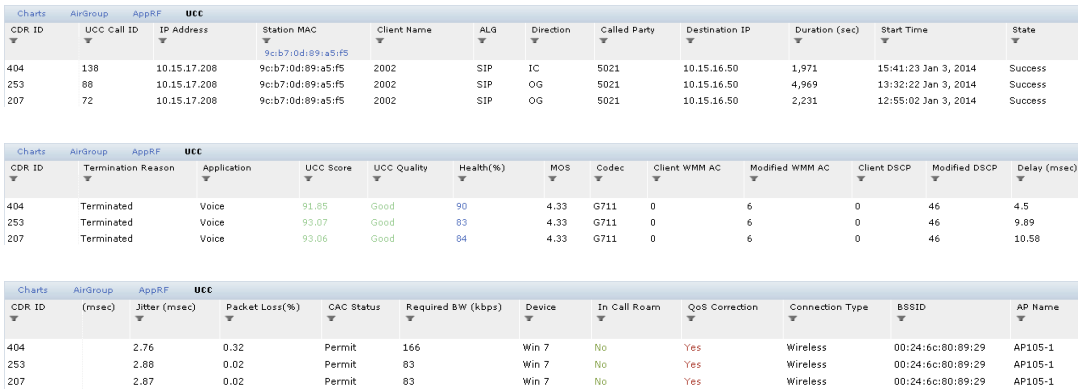


Figure 200 displays all the VoIP call statistics made by a particular client.

Figure 200 *All Calls*



Figure 201 displays the VoIP call summary for a selected call of a client.

Figure 201 *Selected Call Summary*

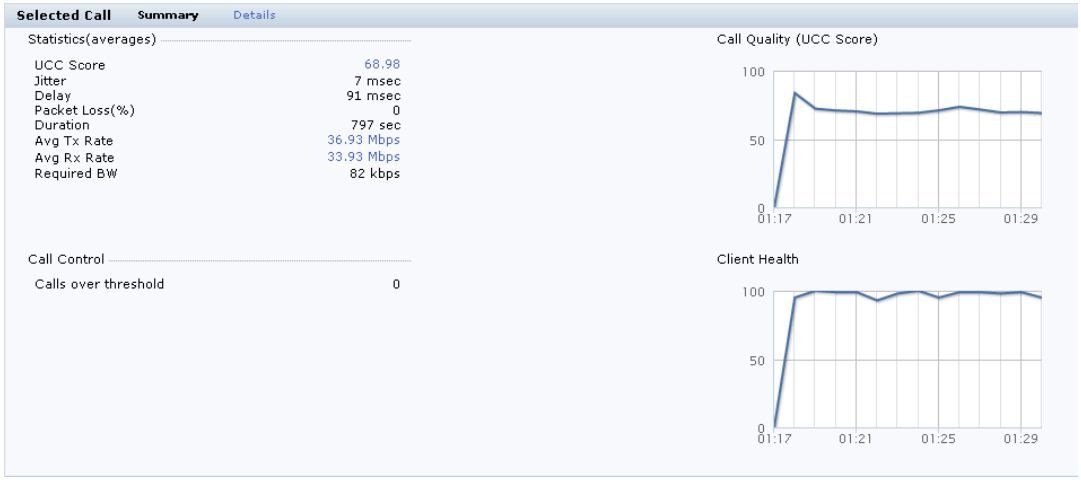
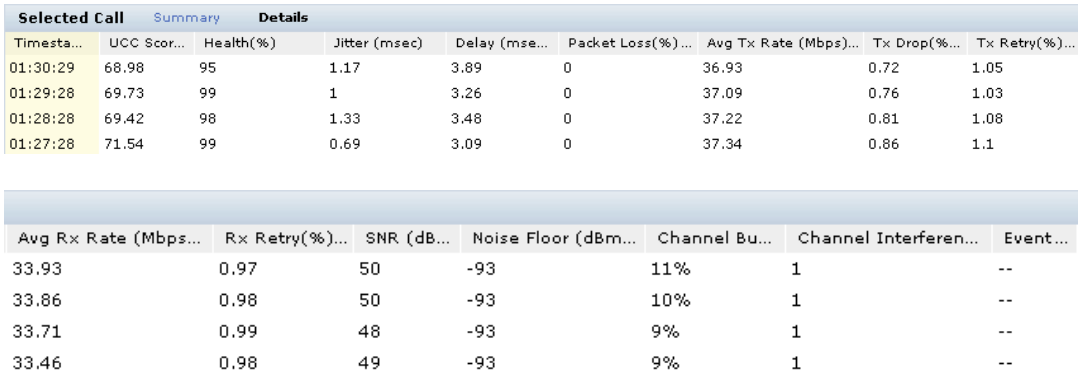


Figure 202 displays the VoIP call details for a selected call of a client.

Figure 202 *Selected Call Details*



Under the **Dashboard > Usage** page of the WebUI, the **Call Quality vs. Client Health** graph displays the co-relation between the VoIP call quality and the VoIP client health of every UCC call. See Figure 203.

Figure 203 *Call Quality vs. Client Health*



Viewing UCC Information

This section describes the procedures to view UCC clients, calls, and configuration information in the controller.



For detailed command parameters, see the *ArubaOS 6.4 CLI Reference Guide*.

- Viewing the list of Lync Clients on page 906
- Viewing UCC Client Information on page 912
- Viewing UCC Configuration on page 912
- Viewing UCC Statistics on page 912
- Viewing UCC Trace Buffer on page 912

Viewing UCC Call Detailed Record

Use the following command to display the CDR statistics for UCC:

```
(host) #show ucc call-info cdrs [ap | app | cid | detail]
```

Viewing UCC Client Information

Use the following command to display the UCC client status and CDR statistics:

```
(host) #show ucc client-info [app | detail | sta]
```

Viewing UCC Configuration

Use the following command to display the UCC configuration in the controller:

```
(host) #show ucc configuration [cac-alg | dialplan | logging | midcall-timeout | realtime-
analysis | rtcp-inactivity | sip | traffic-control]
```

Viewing UCC Statistics

Use the following command to display the UCC call statistics in the controller:

```
(host) #show ucc statistics {counter | dial-plan | remote | tspec-enforcement | wmm-flow}
```

Viewing UCC Trace Buffer

Use the following command to display the UCC call message trace buffer for Lync, SCCP, and SIP ALGs. Events such as establishing voice, video, desktop sharing, and file transfer are recorded:

```
(host) #show ucc trace-buffer {lync | sccp | sip}
```

UCC Troubleshooting

Use the following commands to collect debug information for UCC:

```
(host) (config) #logging level debugging user process stm subcat voice
(host) (config) #logging level debugging user process ucm
(host) (config) #logging level debugging system process ucm
```

UCC-AirWave Integration

The UCC-AirWave integration provides a multi-controller visibility into the UCC solution across deployments. The controller sends raw UCC data using Application MONitoring (AMON) periodically. AirWave Management Platform (AMP) receives these AMON messages and uses this data to display user-friendly aggregated and per-client UCC statistics in AirWave. This helps the administrator to assess the overall health and troubleshoot UCC deployments in a multi-controller environment. The UCC dashboard is supported in AirWave 8.0 onwards.

To register the AMP server with the controller, use the following command to enter the IP address of the AMP server that should receive the AMON messages from the controller:

```
(host) (config) #mgmt-server type amp primary-server <primary-server-ip>
```

UCC Call Quality Metrics

Computing the call quality metrics is an important aspect of troubleshooting. The metrics enable administrators to get an idea of the quality of service on the network and troubleshoot network congestion whenever required. ArubaOS 6.3.x provides UCC call quality metrics for Lync voice calls only at the end of the call. In addition to capturing call quality metrics at the end of the call, ArubaOS 6.4 captures call quality metrics for active calls. The call quality metrics are extended to voice, video, desktop-sharing, and file-transfer applications.

A new metric, UCC score is introduced in ArubaOS 6.4. UCC score computes the quality of voice calls. It takes delay, jitter, and packet loss of RTP packets into account. UCC score is computed on a scale of 0 to 100. This feature works in tunnel, decrypt-tunnel, and split-tunnel forward modes. To compute the UCC score, you must enable RTP Analysis on the master controller. Issue the following CLI commands:

```
(host) (config) #voice real-time-config
(host) (Configure Real-Time Analysis) #config-enable
```

Table 197 shows the call quality parameters displayed on the controller for various UCC ALGs.

Table 197: *Voice and Video Call Quality Parameters*

UCC Application	Media Type	Call Quality Parameters
Lync	Audio	• Mean Opinion Score (MOS) • UCC Score • Delay • Jitter • Packet Loss
	Video	• Delay • Jitter • Packet Loss
SIP	Audio	• UCC Score • Delay • Jitter • Packet Loss
	Video	• Delay • Jitter • Packet Loss
SCCP	Audio	• UCC Score • Delay • Jitter • Packet Loss
NOE	Audio	• UCC Score • Delay • Jitter • Packet Loss
Vocera	Audio	• UCC Score • Delay • Jitter

UCC Application	Media Type	Call Quality Parameters
		• Packet Loss
H.323	Audio	• UCC Score • Delay • Jitter • Packet Loss
	Video	• Delay • Jitter • Packet Loss

Table 198 shows the quality parameters displayed for Lync collaborative services.

Table 198: *Quality Parameters for Collaborative Services*

Lync Collaborative Services	Quality Parameters
Lync Desktop-sharing	• UCC Score • Quality band • Delay (msec) • Jitter (msec) • Signal to Noise Ratio (SNR) • Avg Tx Rate (Mbps) • Tx Drop (%) • Tx Retry (%) • Avg Rx Rate (Mbps) • Rx Retry (%)
Lync File-transfer	• SNR • Avg Tx Rate (Mbps) • Tx Drop (%) • Tx Retry (%) • Avg Rx Rate (Mbps) • Rx Retry (%) **NOTE:** The quality parameters are computed by the AP and does not have any dependency on the quality update message from the Microsoft Lync server.

Changes to Call Admission Control

In ArubaOS 6.4,CAC is not applied for Lync calls. Lync calls are allowed to flow with high priority irrespective of the call count or bandwidth based CAC limit. This applies to calls prioritized by both media classification and

SDN API based Lync ALG. CAC configured under **wlan voip-cac-profile** and Virtual AP (VAP) based bandwidth limitation under **wlan traffic-management-profile** does not apply to Lync calls.

UCC Limitations

- Voice ALGs should not be enabled when voice clients are behind a NAT.
- Media classification does not work when user VLAN has IP NAT configured.
- UCC score is calculated for voice calls and desktop-sharing sessions only.
- For Lync calls, MOS is generated only for voice calls. Lync server does not generate MOS for video calls, desktop-sharing, and file-transfer.

Understanding Extended Voice and Video Features

This section describes the other voice and video-related functionalities that are available on the controller.

Understanding QoS for Microsoft Lync and Apple Facetime

Voice and video devices use a signaling protocol to establish, control, and terminate voice and video calls. These control or signaling sessions are usually permitted using pre-defined ACLs. If, however, the control signaling packets are encrypted, the controller cannot determine which dynamic ports are used for voice or video traffic. In these cases, the controller has to use an ACL with the **classify-media** option enabled to identify the voice or video flow based on a deep packet inspection and analysis of the actual traffic.

Microsoft Lync

Microsoft Lync uses SIPS to establish, control, and terminate voice and video calls. The following example creates an ACL named **lync acl** for Microsoft Lync traffic that identifies port 5061 as the reserved SIP-TLS port.

```
(host) (config) #ip access-list session lync-acl
(host) (config-sess-lync-acl)#any any tcp 5061 permit position 1 queue high classify-media
(host) (config-sess-lync-acl)#any any udp 1025-65535  permit  position 2 queue low
```



You must disable Lync ALG if you use the **classify-media** option. More more information on Lync ALG, see Lync ALG.



Use either Lync ALG or Media Classification because both provide the same functionality.



Media classification is not supported when clients are accessed through a network address translation (NAT).

Microsoft Lync Support for Mobile Devices

Microsoft Lync supports the mobile devices that are running on the following operating systems:

- Windows
- Android
- iOS

You can configure the following ACLs to support the media classification:

- TCP Port 5061 for SIPS signaling sessions initiated by the Lync clients
- TCP Port 443 for signaling sessions initiated by Lync application running on mobile devices
- UDP and TCP traffic on port range 1024 to 65535 for sessions initiated by the Lync applications

The following example shows how to configure an ACL to identify and monitor the mobile devices supported by Lync:

```
(host) (config) #ip access-list session Lync-Smart-Device
(host) (config-sess-Lync-Smart-Device)#any ?
(host) (config-sess-Lync-Smart-Device)#any alias Lync-Servers ?
(host) (config-sess-Lync-Smart-Device)#any alias Lync-Servers tcp 443 ?
(host) (config-sess-Lync-Smart-Device)#any alias Lync-Servers tcp 443 permit classify-media
(host) (config-sess-Lync-Smart-Device)#any any ?
(host) (config-sess-Lync-Smart-Device)#any any udp 1025-65535  permit  position 3 queue low
```

Apple Facetime

When an Apple device starts a Facetime video call, it initiates a TCP session to the Apple Facetime server over port 5223, then sends SIP signaling messages over a non-default port. When media traffic starts flowing, audio and video data are sent through that same port using RTP. (The audio and video packets are interleaved in the air, though individual sessions can be uniquely identified using their payload type and sequence numbers.) The RTP header and payload also get encapsulated under the TURN ChannelData Messages. The Facetime call is terminated with a SIP BYE message that can be sent by either party.

Table 199 lists the ports used by Apple Facetime. Facetime users need to be assigned a role where traffic is allowed on these ports:

Table 199: *Ports used by the Apple Facetime Application*

Port	Packet Type
53	TCP/UDP
443	TCP
3478-3497	UDP
5223	TCP
16384-16387	UDP
16393-16402	UDP

The example below shows how to configure an ACL to identify and monitor Apple Facetime traffic:

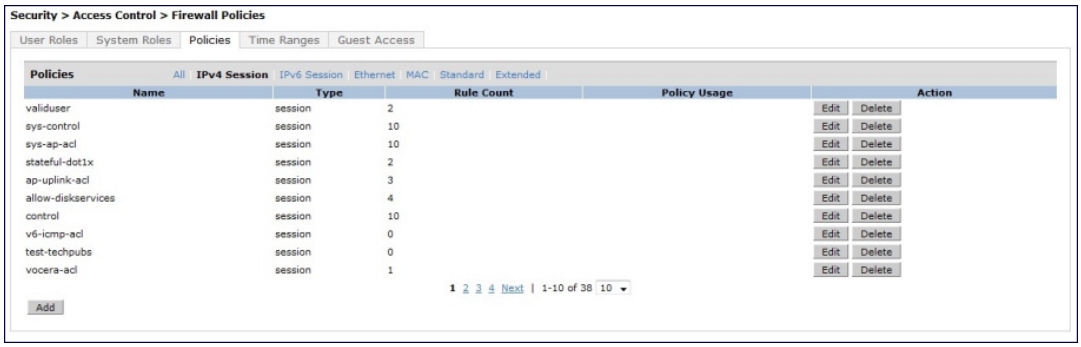
```
(host) (config) #ip access-list session facetime-acl
(host) (config-sess-facetime-acl)#any any tcp 80 permit position 1 queue low
(host) (config-sess-facetime-acl)#any any tcp 443 permit position 2 queue low
(host) (config-sess-facetime-acl)#any network 17.0.0.0 255.0.0.0 tcp 5223 permit position 3
queue low classify-media
(host) (config-sess-facetime-acl)#any any UDP 80 permit position 4 queue low
(host) (config-sess-facetime-acl)#any network 17.0.0.0 255.0.0.0 UDP 16384-16387 permit
position 5 queue low
```

You can use the WebUI or CLI for enabling the **Classify Media** option for the encrypted signaling protocols. In our example, we will configure this support for Microsoft Lync:

In the WebUI

1. Navigate to the **Configuration > Security > Access control** page.
2. Click the **Policies** tab.

Figure 204 *Firewall Policies Tab*

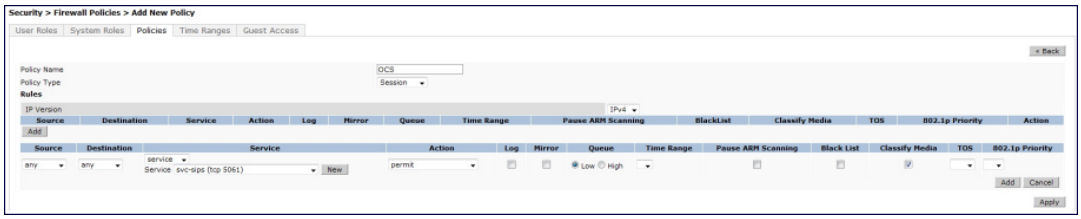


3. Click **Add** to create a new policy.
4. Enter a name for the policy in the **Policy Name** field and choose **Session** in the **Policy Type** drop down menu.
5. Select **IPv4** in the **IP Version** drop down menu and click **Add**.
6. In the **Service** column, choose **service** and **Select svc-sips (tcp-5061)** from the **Service** drop-down menu.
7. Select the **Classify Media** check box.



There will be a performance impact if you choose **any** in the **Service** column and enable the **Classify Media** flag for the deep packet inspection.

Figure 205 *Enabling Classify Media*



8. Click **Apply**.

Enabling WPA Fast Handover

In the 802.1x Authentication profile, the WPA fast handover feature allows certain WPA clients to use a pre-authorized PMK, significantly reducing handover interruption. Check with the manufacturer of your handset to see if this feature is supported. This feature is disabled by default.



This feature supports WPA clients, while opportunistic key caching (also configured in the 802.1x Authentication profile) supports WPA2 clients.

In the WebUI

1. Navigate to the **Configuration > AP Configuration** page. Select either **AP Group** or **AP Specific**.
   - If you select **AP Group**, click **Edit** for the AP group name for which you want to enable WPA fast handover.
   - If you select **AP Specific**, select the name of the AP for which you want to enable WPA fast handover.
2. Under **Profiles**, select **Wireless LAN**, then **Virtual AP**. In the **Virtual AP** list, select the appropriate virtual AP instance.

3. Select **AAA** profile. Select the **802.1x Authentication Profile** to display in the **Profile Details** section.
4. Scroll down to select the **WPA-Fast-Handover** check box.
5. Click **Apply**.

In the CLI

Use the following commands:

```
aaa authentication dot1x <profile>
   wpa-fast-handover
```

For deployments where there are expected to be considerable delays between the controller and APs (for example, in a remote location where an AP is not in range of another Dell AP) you can increase the value for the bootstrap threshold in the AP System profile to minimize the chance of the AP rebooting due to temporary loss of connectivity with the Dell controller.

Enabling Mobile IP Home Agent Assignment

When you enable IP mobility in a mobility domain, the proxy mobile IP module determines the home agent for a roaming client. An option related to voice clients that you can enable allows on-hook phones to be assigned a new home agent to load balance voice client home agents across controllers in the mobility domain. See IP Mobility on page 599 for more information about mobility.

Scanning for VoIP-Aware ARM

ARM scanning on an AP during a call affects the voice quality. You can pause the ARM scanning on the AP when a call is active by turning on the VoIP-Aware ARM Scanning support to avoid voice quality issues.

You can use the WebUI or CLI to enable VoIP-aware ARM scanning in the ARM profile.

In the WebUI

1. Navigate to the **Configuration > AP Configuration** page. Select either the **AP Group** or **AP Specific**.
   - If you selected the **AP Group** tab, click **Edit** by the name of the AP group with the ARM profile you want to configure.
   - If you selected the **AP Specific** tab, click **Edit** by the name of the AP with the ARM profile you want to configure.
2. In the **Profiles** list, Expand the **RF Management** section.
3. Select **Adaptive Radio Management (ARM) Profile**.
4. Select a profile instance from the drop-down menu to edit that profile.
5. Select the **VoIP Aware Scan** option.
6. Click **Apply**.

For additional information on configuring an Adaptive Radio Management profile, see Configuring ARM Profiles on page 435.

In the CLI

```
rf arm-profile <profile-name>
   voip-aware-scan
```

Disabling Voice-Aware 802.1x



The Voice-Aware 802.1x support is deprecated for ArubaOS 5.0 and later releases.

Although reauthentication and rekey timers are configurable on a per-SSID basis, an 802.1x transaction during a call can affect voice quality. If a client is on a call, 802.1x reauthentication and rekey are disabled by default until the call is completed. You disable or re-enable the "voice aware" feature in the 802.1x authentication profile.

In the WebUI

1. Navigate to the **Configuration > AP Configuration** page. Select either AP Group or AP Specific.
   - If you select **AP Group**, click **Edit** for the AP group name for which you want to disable voice awareness for 802.1x.
   - If you select **AP Specific**, select the name of the AP for which you want to disable voice awareness for 802.1x.
2. Under Profiles, select **Wireless LAN**, then select **Virtual AP**. In the Virtual AP list, select the appropriate virtual AP instance.
3. Select **AAA profile**. Select the **802.1x Authentication Profile** to display in the **Profile Details** section.
4. Scroll down and deselect the **Disable rekey and reauthentication for clients on call** check box.
5. Click **Apply**.

In the CLI

Use the following commands:

```
aaa authentication dot1x <profile>
   no voice-aware
```

Configuring SIP Authentication Tracking

The controller supports the stateful tracking of session initiation protocol (SIP) authentication between a SIP client and a SIP registry server. Upon successful registration, a user role is assigned to the SIP client. You specify a configured user role for the SIP client in the AAA profile.

In the WebUI

1. Navigate to the **Configuration > AP Configuration** page. Select either **AP Group** or **AP Specific**.
   - If you select **AP Group**, click **Edit** for the AP group name for which you want to configure the SIP client user role.
   - If you select **AP Specific**, select the name of the AP for which you want to configure the SIP client user role.
2. Under **Profiles**, select Wireless LAN, then **Virtual AP**. In the **Virtual AP** list, select the appropriate virtual AP instance.
3. Select the AAA profile. Enter the configured user role for SIP authentication role.
4. Click **Apply**.

In the CLI

Use the following commands:

```
aaa profile <profile>
   sip-authentication-role <role>
```

Use the **show voice client-status** command to view the state of the client registration.

Enabling Real Time Call Quality Analysis

Real Time Call Quality Analysis (RTCQA) enables the controller to compute the call quality parameters such as jitter, delay, packet loss, and call quality score (R-value) directly from the RTP media stream. Additionally, the

controller saves the periodic samples of the quality parameters for detailed analysis of the results. You can monitor up to 30 active calls that are initiated after enabling RTCQA. You can avail the full benefits of RTCQA by setting the AP in tunnel, decrypt-tunnel, or split-tunnel forwarding mode.

Enabling RTCQA is helpful in cases where the VOIP clients do not use RTP Control Protocol (RTCP) or use encrypted RTCP (in the case of Lync) which the controller cannot get the quality information from the RTCP frames.

Important Points to Remember

RTCQA for voice calls is supported only in the following cases:

- when the signaling messages are not encrypted.
- when the signaling messages are encrypted for Lync.
- when the voice client does not roam from one controller to another controller. In other words, when a client moves to a foreign agent controller, RTCQA does not take effect.

You can use the WebUI or CLI to enable RTCQA and view the call quality reports based on the analysis.

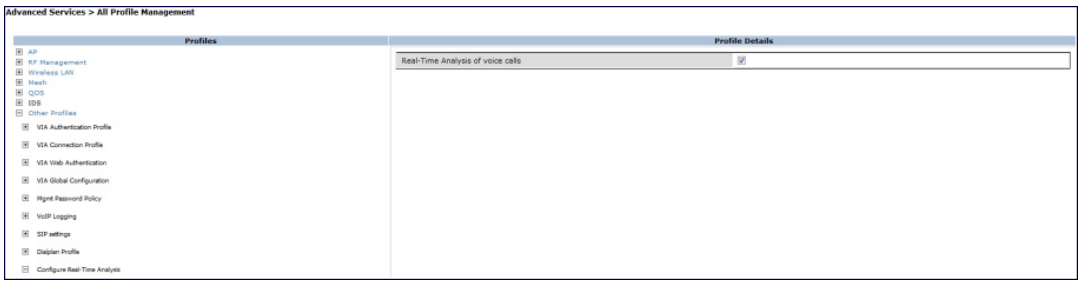
NOTE

To generate Lync UCC score, you must enable RTCQA. For more information on UCC score, see UCC Call Quality Metrics on page 913

In the Web UI

1. Navigate to the **Configuration > Advanced services > All Profiles** page.
2. Expand **Other Profiles** under the **Profiles** section and click **Configure Real-Time Analysis**.
3. Enable Real Time call quality analysis for the voice calls by selecting the **Real-Time Analysis of voice calls** check box.

Figure 206 *Enable Real Time Analysis*



4. Click **Apply**.

Viewing Real Time Call Quality Reports

1. To view the average Real Time analysis reports, navigate to the **Monitoring > Voice > Real-Time Quality Analysis** page.
2. To view the detailed Real Time analysis report of a specific client, select the client and click **View Details**.

NOTE

Real Time analysis report is not available for clients in tunnel or bridge mode.

In the CLI

To configure Real Time analysis on voice calls:

```
(host) (config) #voice real-time-config
(host) (Configure Real-Time Analysis) #config-enable
```

To view the average Real Time analysis reports for the voice clients:

```
(host) #show voice real-time-analysis

Real-Time Analysis Call Quality Report
--------------------------------------
Client(IP)     Client(MAC)         Client(Name)   ALG   Jitter(D)(usec)   Pkt-loss(D)(%)
----------     -----------         ------------   ---   ---------------   --------------
10.16.33.251   00:1f:6c:7a:d4:fd   6005           sccp  16.980            0.625
10.15.16.201   1c:ab:a7:2d:75:6b   7129           Lync  10.934            0.120

Delay(D)(usec)    UCC Score(D)   Forward mode
--------------    ------------   ------------
421.125           78.985         decrypt-tunnel
454.163           86.250         tunnel

Num Records:2
```

To view the detailed Real Time analysis report for a specific client:

```
(host) #show voice real-time-analysis sta 1c:ab:a7:2d:75:6b
WARNING: This command will be deprecated, please use show ucc commands instead

Real-Time Analysis Detailed Report
----------------------------------
Time               Jitter(D)(usec)   Pkt-loss(D)(%)   Delay(D)(usec)   UCC Score(D)   Forward mode
----------------   ---------------   --------------   --------------   ---------      ------------
Mar 15 17:05:34    2.000             1.000            255.000          88.360         tunnel
Mar 15 17:05:32    2.000             5.000            211.000          78.360         tunnel
Mar 15 17:05:30    3.000             7.000            203.000          73.360         tunnel
Mar 15 17:05:28    2.000             2.000            271.000          86.360         tunnel
```

Enabling SIP Session Timer

SIP session timer is implemented in the SIP ALG as per RFC 4028.

SIP session timer defines a keep alive mechanism for the SIP sessions using the periodic session refresh requests from the user agents. The interval for the session refresh requests is determined through a negotiation mechanism. If a session refresh request is not received within the negotiated interval, the session is assumed to be terminated.

For more information on the SIP session timer support, See *section 8.0, Proxy Behaviour* in the RFC 4028.



This release of ArubaOS does not support the configurable Min-SE parameter for SIP ALG. Therefore, the ALG will not generate the 422 responses for the session refresh requests.

You can use the WebUI or CLI to enable the SIP session timer and set the session-expiry timer value using the WebUI and CLI.

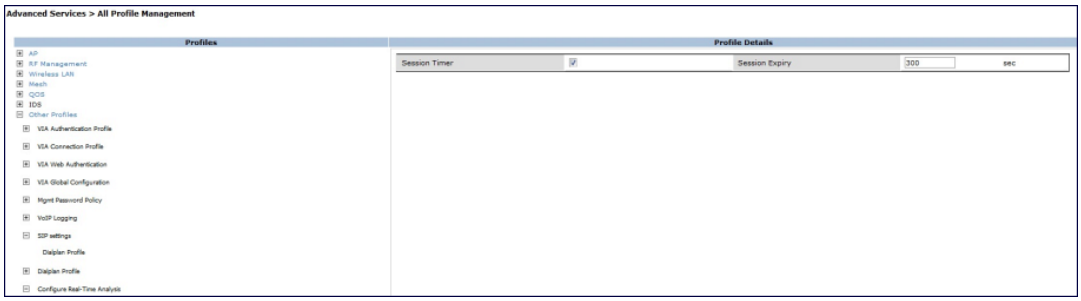


SIP Session Timer can be configured only for SIP over UDP.

In the WebUI

1. Navigate to the **Configuration > Advanced services > All Profiles** page.
2. Expand **Other profiles** under the **Profiles** section and click **SIP Settings**.
3. Enable the session timer by selecting the **Session Timer** check box under the **Profile Details** section.
4. Specify a timeout value in seconds in the **Session Expiry** field. The range is 240- 1200 seconds. The default value is 300 seconds.

Figure 207 *Enabling SIP Session Timer*



5. Click **Apply**.

In the CLI

To configure the session timer and the timeout value:

```
(host) #configure terminal
(host) (config) #voice sip
(host) (SIP settings) #session-timer
(host) (SIP settings) #session-expiry 400
```

To view the SIP settings on the controller:

```
(host) #show voice sip

SIP settings
------------
Parameter          Value
---------          -----
Session Timer      Enabled
Session Expiry     400 sec
Dialplan Profile   N/A
```

Enabling Wi-Fi Edge Detection and Handover for Voice Clients

Voice clients in an infrastructure can be switched to an alternate carrier or connection when they leave their active Wi-Fi coverage or roam to an area with poor Wi-Fi coverage. The controller uses the best Wi-Fi signal strength (dbm value) reported by the voice clients (received from all APs) to determine if the voice clients are within or leaving their active Wi-Fi connection. If the signal strength is weak, the controller will trigger the handover process to switch the voice client to an alternate carrier or connection. This process ensures QoS for voice calls.



- The handover process is available for voice clients supporting the 802.11K standard and with the ability to transmit and receive beacon reports.
- The voice clients should have dual mode capabilities to ensure that they can switch to an alternate network in case of a loss in Wi-Fi coverage.

The handover process can be configured using the **wlan handover-trigger-profile** command. Use the **handover-threshold** parameter to specify the threshold value (dbm) and enable the **handover-trigger** parameter. If the best signal strength reported by a voice client is equal to or less than the threshold value, the handover process is initiated.

In the WebUI

1. Navigate to the **Configuration > Advanced Services > All Profiles** page.
2. Expand **Wireless Lan** under the **Profiles** section.

3. Expand **802.11 K** profile under **Wireless Lan**.
4. Select the default profile.
5. Select **Advertise 802.1k Capability**.
6. In the **profiles** list, note which Handover Trigger Feature Settings profile is associated with the selected 802l11k profile.
7. Expand **Handover Trigger** under **Wireless Lan**.
8. Select the handover trigger profile associated with the default 802.11k profile.
9. Select the **Enable Handover Trigger feature** checkbox
10. Specify the handover threshold value in the **Threshold signal strength value at which handover Trigger should be sent to the client** field. The handover threshold value should be within the range 20 to 70 dbm. The default threshold value is -60 dbm.
11. Click **Apply**.

In the CLI

The following command enables the dot11k profile and sets the handover threshold at -60dbm.

```
(host) (config) #wlan handover-profile default
(host) (802.11K Profile "default") #dot11k-enable
(host) (802.11K Profile "default") #handover-trigger-profile default
(host) (802.11K Profile "default") #exit
(host) (config) #wlan handover-trigger-profile default
(host) (Handover Trigger Profile) #handover-trigger
(host) (Handover Trigger Profile) #handover-threshold 60
```



The handover threshold value is a negative dbm value. In the CLI, enter the value without the negative (-) sign.

Working with Dial Plan for SIP Calls

A PSTN call from a SIP device usually requires the user to prefix 9 or 0 before the destination number. You can configure dial plans (prefix codes) on the controller that are required by the local EPABX system to provide outgoing PSTN call facility from a SIP device. After the dial plan is configured, a user can make SIP calls by dialing the destination number without any prefixes.



Dial plan can be configured only for SIP over UDP.

Understanding Dial Plan Format

The format of a SIP dial plan is `<sequence> <pattern> <action>`.

- sequence—is a number between 100 and 65535. The sequence number positions the dial plan in the list of dial plans configured in the controller.
- pattern—is the digit pattern or the number of digits that will be dialed by the user. You can specify digit pattern using 'X', 'Z', 'N', '[]', and '.'.
  - X is a wild card that represents any character from 0 to 9.
  - Z is a wild card that represents any character from 1 to 9.
  - N is a wild card that represents any character from 2 to 9.
  - . (period) is a wild card that represents any-length digit strings.
- action—is the prefix code that is automatically prefixed to the dialed number. This is specified as **<prefix-code>%e**. Examples of prefix codes are:

- 9%e: The number 9 is prefixed to the dialed number.
- 91%e: The number 91 is prefixed to the dialed number.

Table 200: *Examples of Dial Plans*

Dialplan Pattern	Action	Description
XXXX	%e	When the user dials a four digit number, no action is taken and the call is allowed.
XXXXXXX	9%e	When the user dials a seven digit number, a nine (**9**) is prefixed to that number and the call is executed. Example, if the user dials 2274500, the call is executed by adding 9 to the number, 92274500.
XXXXXXXXXX	91%e	This dial plan prefixes **91** to the dialed number. Example, call to 4082274500 will be executed as 914082274500.
+1XXXXXXXXXX	9%e	This dial plan replaces '**+**' with **9** and executes the call. Example, call to +14082274500 is executed as 914082274500.
+.	9011%e	This dial plan removes '+' and prefixes 9011 for an international call. Example, call to +886212345678 is executed as 9011886212345678.

Configuring Dial Plans

You can configure a maximum of two dial plan profiles and maximum of 20 dial plans per profile. The dial plan must be associated to a SIP ALG configuration.

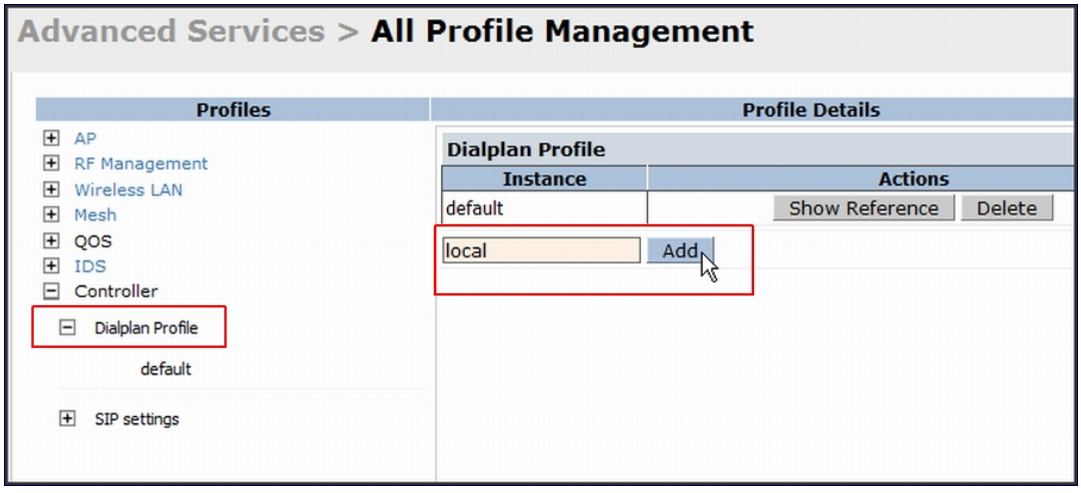
To configure a dial plan for SIP devices:

1. Create a voice dial plan.
2. Associate the dial plan with SIP ALG.

In the WebUI

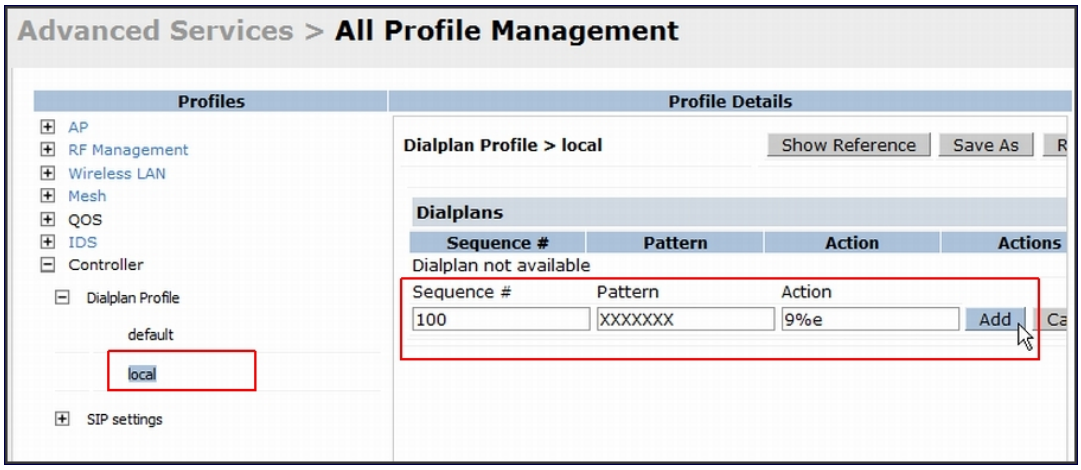
1. In the WebUI, navigate to **Configuration > Advanced Services > All Profiles > Controller > Dialplan Profile**. Enter a name for the dial plan profile and click **Add**.

Figure 208 *Dialplan Profile*



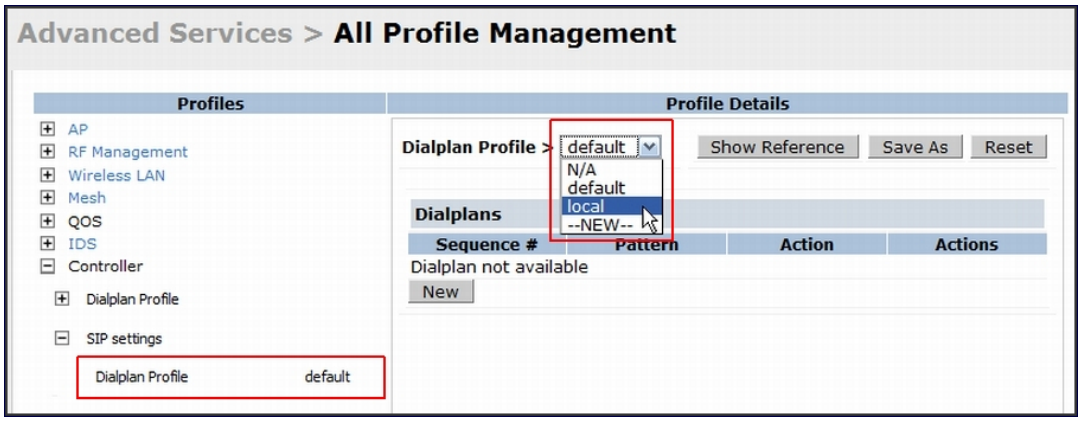
2. Under **Profiles**, expand **Controller** and select the newly created dial plan profile. Enter the following dial plan details and click **Add**.
   - Sequence number: the dial plan position in the list of dial plans
   - Pattern: the number that the user will dial
   - Action: prefix to be added by the **controller** before forwarding the call to the EPABX

Figure 209 *Dialplan Details*



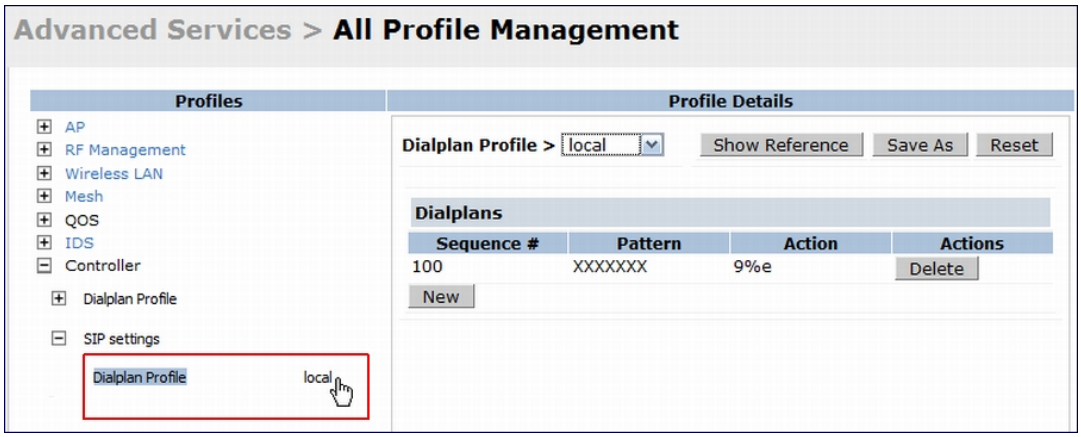
3. Click **Apply**.
4. Under **Profile**, navigate to **Controller > SIP settings** and select **Dialplan Profile.** In the **Profile Details section,** select the **Dialplan Profile** from the drop-down list and click **Apply** .

Figure 210 ***Select Dialplan Profile***



The Dialplan Profile displays the dial plan details:

Figure 211 *View Dialplan Details*



In the CLI

To create a voice dial plan profile:

```
(host) (config) #voice dialplan-profile local
(host) (Dialplan Profile "local") #dialplan 100 XXXXXXX 9%e
(host) (Dialplan Profile "local") #!
```

To associate the dial plan with SIP ALG:

```
(host) (config) #voice sip
(host) (SIP settings) #dialplan-profile local
(host) (SIP settings) #!
```

To view the SIP dial plan profile:

```
(host) (config) #show voice sip

SIP settings
------------
Parameter         Value
---------         -----
Dialplan Profile  local
```

To view the dial plan details:

```
(host) (config) #show voice dialplan-profile local

Dialplan Profile "local"
--------------------------
Parameter  Value
---------  -----
dialplan   100 XXXXXXX 9%e
```

Enabling Enhanced 911 Support

ArubaOS provides seamless support for emergency calls in the Dell network by interoperating with RedSky emergency call server. The controller uses SNMP to interoperate with RedSky call handling system.



This release of ArubaOS supports only RedSky emergency call server.

You must configure the RedSky server as an SNMP host and enable SNMP traps to activate the E911 feature on the controller. For more information on configuring the RedSky server as SNMP host, see Configuring SNMP on page 801.

The E911 support has the following basic functions:

- location tracking
- call handling
- caller identification and callback capability

For information on call handling, caller identification and callback capability, see the RedSky documentation.

The controller tracks the location of the voice clients and notifies the emergency call server using SNMP traps. The controller notifies the location of a voice client to the emergency server:

- when it identifies a voice client
- when a voice client roams from one access point to another access point in the same controller
- when a voice client roams from one access point to another access point in a different controller
- when a voice client registers with a PBX system

The notification process ensures that the emergency call server is notified whenever a voice client is identified or the location of the client is updated. If a voice client roams outside of a WLAN coverage, the controller does not send any notifications to the emergency call handling system. This may happen when there is a sudden loss of WLAN coverage due to extreme conditions such as fire accident. In such cases, the last associated access point will be the location of the voice client.



The controller tracks the location only for voice clients. To track the location of a remote voice client, the administrator must configure the location of the remote access point in the controller or emergency call server.

The emergency call server queries the controller using the SNMP :get: request to get the location of a specific emergency caller. In response to the location query, the controller sends the following parameters to the emergency server:

- Client IP Address
- Client Mac Address
- AP Name
- AP Wired MAC
- AP Location
- AP Mode
- Controller IP Address

The controller also supports location queries for the clients that are not identified as voice clients on the controller.

Working with Voice over Remote Access Point

Voice traffic support is enhanced on split tunnel mode over a remote access point. The voice traffic management for remote and local users are done on the controller. However, the sessions are created differently for both users. For remote users, the sessions are created on the remote access point and for local users, the sessions are created on the controller. This enhancement provides the following support for the voice traffic in the split tunnel over remote access point:

- voice traffic QoS is consistent for both local and remote users
- all voice ALGs work reliably in split tunnel mode when the PBX traffic is destined to flow through the corporate network.
- provides voice statistics and counters for remote voice clients in the split tunnel mode

  The `flag` parameter in the `show voice client-status` command is updated to indicate remote users:

```
(host) #show voice client-status
```

```
Voice Client(s) Status
----------------------
AP Name  BSSID                ESSID Client(MAC)    Client(IP)    Client Name  Server(IP)
Registration State   Call Status  ALG   Flags

-------  -----                -----  -----------    ----------    -----------  ----------  --
---------------  ----------  ---   -----

moscato 00:0b:11:5c:d6:80 home   00:00:5c:04:b3:10   10.20.1.100   Client        10.13.8.1
REGISTERED          Idle        h323  R

Num Clients:1
Flags:    R - Remote user
```

Understanding Battery Boost

Battery boost is an optional feature that can be enabled for any SSIDs that support voice traffic. This feature converts all broadcast and multicast traffic to unicast before delivery to the client. Enabling battery boost on an SSID allows you to set the DTIM interval from 10 to 100 (the previous allowed values were 1 or 2), equating to 1,000 to 10,000 milliseconds. This longer interval keeps associated wireless clients from activating their radios for multicast indication and delivery, leaving them in power-save mode longer, and thus lengthening battery life. The DTIM configuration is performed on the WLAN, so no configuration is necessary on the client.

An associated parameter available on some clients is the Listening Interval (LI). This defines the interval (in number of beacons) after which the client must wake to read the Traffic Indication Map (TIM). The TIM indicates whether there is buffered unicast traffic for each sleeping client. With battery boost enabled, the DTIM is increased but multicast traffic is buffered and delivered as unicast. Increasing the LI can further increase battery life, but can also decrease client responsiveness.



Do not enable battery boost if your network includes Polycom SpectraLink devices that use the Push-to-Talk feature.

You can use the WebUI or CLI to enable the battery boost feature and set the DTIM interval in the SSID profile.

In the WebUI

1. Navigate to the **Configuration > AP Configuration** page. Select either the **AP Group** tab or **AP Specific** tab.
   - If you selected **AP Group**, click **Edit** by the AP group name for which you want to enable battery boost.
   - If you selected **AP Specific**, select the name of the AP for which you want to enable battery boost.
2. Under **Profiles**, select **Wireless LAN**, then **Virtual AP**. In the **Virtual AP** list, select the appropriate virtual AP instance.
3. In the **Profile Details** section, select the SSID profile you want to configure.
4. Click the **Advanced** tab.
5. Scroll down the Advanced options and select the **Battery Boost** check box.
6. Scroll up to change the **DTIM** Interval to a longer interval time.
7. Click **Apply**.

In the CLI

Use the following commands:

```
wlan ssid-profile <profile>
   battery-boost
   dtim-period <milliseconds>
```

Enabling LLDP

Link Layer Discovery Protocol (LLDP), is a Layer-2 protocol that allows network devices to advertise their identity and capabilities on a LAN. Wired interfaces on Dell APs support LLDP by periodically transmitting LLDP Protocol Data Units (PDUs) comprised of selected type-length-value (TLV) elements. For a complete list of supported, see Table 201 and Table 202.

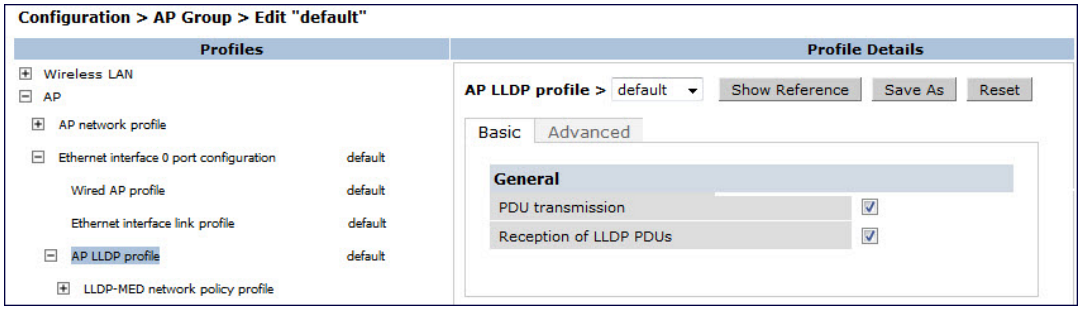
LLDP-MED (media endpoint devices) is an extension to LLDP that supports interoperability between VoIP and video streaming devices and other networking clients. LLDP-MED network policy discovery lets end-points and network devices advertise the VLAN, priority levels, and DSCP values used by a voice or video application.

In the WebUI

Use the procedure below to configure the LLDP and LLDP-MED profiles and select the TLVs to be sent by the AP.

1. Navigate to the **Configuration > AP Configuration** page. Select either the **AP Group** or **AP Specific** .
   - If you selected **AP Group**, click **Edit** by the AP group name for which you want to enable LLDP.
   - If you selected **AP Specific**, select the name of the AP for which you want to enable LLDP.
2. In the **Profiles** window, expand **AP**, then expand the **Ethernet interface port configuration profile** for the port for which you want to configure LLDP.
3. Select the **AP LLDP Profile**.

Figure 212 *AP LLDP Profile Details*



4. The AP LLDP profile is divided into two tabs, **Basic** and **Advanced**. The **Basic** tab displays only those configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. If you change a setting on one tab, then click and display the other tab without saving your configuration, that setting will revert to its previous value. Both basic and advanced settings are described in Table 171.
5. Configure the LLDP profile parameters as desired then click

Table 201: *LLDP Profile Configuration Parameters*

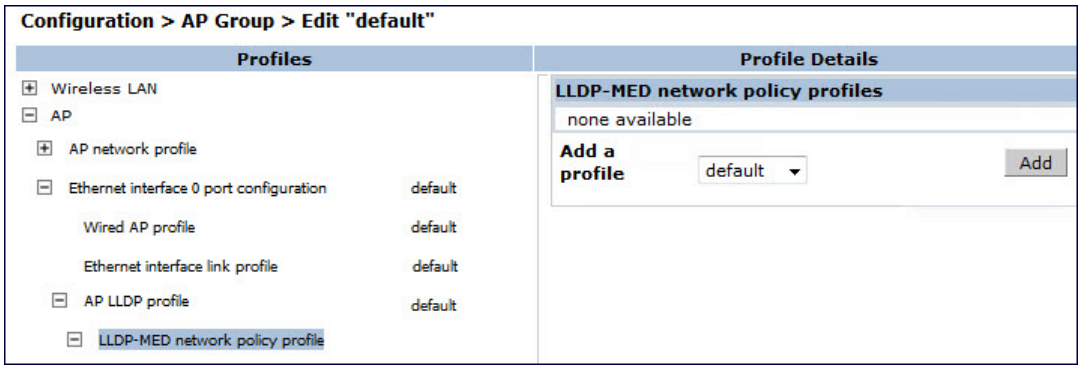
Parameter	Description
Basic Settings	
PDU Transmission	Select this checkbox to enable LLDP PDU Transmission. PDU Transmission is enabled by default.
Reception of LLDP PDUs	Select this checkbox to enable LLDP PDU Reception. PDU Reception is enabled by default.

Parameter	Description
Advanced Settings	
Transmit Interval (seconds)	The interval between LLDP TLV transmission seconds. Range: 1-3600, seconds and Default: 30 seconds.
Transmit hold multiplier	The Transmit hold multiplier is a value that is multiplied by the transmit interval to determine the number of seconds to cache learned LLDP information before that information is cleared. If the Transmit hold multiplier value is set at its default value of 4, and the Transmit interval is at its default value of 30 seconds, then learned LLDP information will be cached for 4x30 seconds, or 120 seconds.
Optional TLVs	Click the checkboxes in this section to select the optional TLVs the AP interface sends in LLDP PDUs. The AP will send all optional TLVs by default. • **port-description:** transmit a TLV that gives a description of the AP's wired port in an alphanumeric format. • **system-description**: transmit a TLV that describes the AP's model number and software version. • **system-name**: transmit a TLV that sends the AP name or wired MAC address. • **capabilities**: transmit the system capabilities TLV to indicate which capabilities are supported by the AP. • **management-address**: transmit a TLV that indicates the AP's management IP address, in either IPv4 or IPV6 format.
802.1 TLVs	Click the checkboxes in this section to select the 802.1 TLVs the AP interface sends in LLDP PDUs. The AP will send all 802.1 TLVs by default: • **port-vlan**: transmit the LLDP 802.1 port VLAN TLV. If the native VLAN is configured on the port, the port-vlan TLV will send that value, otherwise it will send a value of "0". • **vlan-name**: transmit the LLDP 802.1 VLAN name TLV. The AP sends a value of "Unknown" for VLAN 0, or "VLAN <number>" for all non-zero VLAN numbers.
802.3 TLVs	Click the checkboxes in this section to select the 802.3 TLVs the AP interface sends in LLDP PDUs. The AP will send all 803.2 TLVs by default: • **mac**: transmit the 802.3 MAC/PHY Configuration/Status TLV to indicate the AP interface's duplex and bit rate capacity and current duplex and bit rate settings. • **link-aggregation**: transmit the 802.3 link aggregation TLV to indicate that link aggregation is not supported. • **mfs**: transmit the 802.3 Maximum Frame Size (MFS) TLV to show the AP's maximum frame size capability. • **power**:transmit the 802.3 Power Via media dependent interface (MDI) TLV to show the power support capabilities of the AP interface. This parameter is supported by the W-AP130 Seriesonly.
LLDP-MED TLVs	Once you have associated an LLDP-MED Network policy profile with this LLDP profile, you can click the checkboxes in this section to select the LLDP-MED TLVs the AP interface sends in LLDP PDUs. The AP does not send any LLDP-MED TLVs by default: • **capabilities**: transmit the LLDP-MED capabilities TLV. The AP will automatically send this TLV if it sends any other LLDP-MED TLVs. • **inventory**: transmit the LLDP-MED inventory TLV. • **network-policy**: transmit the LLDP-MED network-policy TLV. **NOTE:** The TLVs in this section cannot be enabled unless you have associated an LLDP-MED Network policy profile

6. Click **Apply**.

7. To associate an LLDP-MED network policy profile with the LLDP profile and select the LLDP-MED TLVs to be sent by the AP interface, click the **LLDP-MED network policy profile** that appears below the AP LLDP profile in the profile list:

Figure 213 *AP LLDP Profile Details*



8. If the LLDP profile does not currently reference an LLDP-MED profile, you must associate an LLDP-MED profile with the LLDP profile before you can configure any LLDP-MED settings. Click the **Add a profile** drop-down list in the **Profile Details** window.
   - To associate an existing LLDP-MED network policy, click an LLDP-MED policy name, then click **Add**.
   - To create a new LLDP-MED policy, click NEW, enter a name for the LLDP-MED network policy, then click **Add**.
9. Click **Apply** .
10. Next, expand the LLDP-MED network policy profile in the **Profiles** list, and select the profile you want to configure.
11. The LLDP-MED network policy profile is divided into two tabs, **Basic** and **Advanced**. The **Basic** tab displays only those configuration settings that often need to be adjusted to suit a specific network. The **Advanced** tab shows all configuration settings, including settings that do not need frequent adjustment or should be kept at their default values. If you change a setting on one tab, then click and display the other tab without saving your configuration, that setting will revert to its previous value. Both basic and advanced settings are described in Table 202.
12. Configure the LLDP-MED profile parameters as desired then click

Table 202: *LLDP-MED Profile Configuration Parameters*

Parameter	Description
Basic Settings	
LLDP-MED application type	Click the LLDP-MED application type drop-down list and select the application type managed by this profile. • **guest-voice**: if the AP services a separate voice network for guest users and visitors. • **guest-voice-signaling**: if the AP is part of a network that requires a different policy for guest voice signaling than for guest voice media. Do not use this application type if both the same network policies apply to both guest voice and guest voice signaling traffic. • **softphone-voice**: if the AP supports voice services using softphone software applications on devices such as PCs or laptops. • **streaming-video**: if the AP supports broadcast or multicast video or other streaming video services that require specific network policy treatment. This application type is not recommended for video applications that rely on TCP with buffering. • **video-conferencing**: of the AP supports video conferencing equipment that provides real-time, interactive video/audio services. • **video-signaling**: if the AP is part of a network that requires a different policy for video signaling than for the video media. Do not use this application type if both the same network policies apply to both video and video signaling traffic. • **voice**: if the AP services IP telephones and other appliances that support interactive voice services. This is the default application type. • **voice-signaling**: Select this application type if the AP is part of a network that requires a different policy for voice signaling than for the voice media. Do not use this application type if both the same network policies apply to both voice and voice signaling traffic.
LLDP-MED application VLAN	Specify a VLAN by VLAN ID (0-4094) or VLAN name.
LLDP-MED application VLAN tagging	Click this checkbox if the LLDP-MED policy applies to a to a VLAN that is tagged with a VLAN ID or untagged. The default value is untagged. **NOTE:** When an LLDP-MED network policy is defined for use with an untagged VLAN, then the L2 priority field is ignored and only the DSCP value is used.
Advanced Settings	
LLDP-MED application Layer-2 priority	Specify a 802.1p priority level for the specified application type, by entering a value from 0 to 7, where 0 is the lowest priority level and 7 is the highest priority.
LLDP-MED application Differentiated Services Code Point	Select a Differentiated Services Code Point (DSCP) priority value for the specified application type by specifying a value from 0 to 63, where 0 is the lowest priority level and 63 is the highest priority.

13. Click **Apply**.

In the CLI

Use the following command:

```
ap lldp profile <profile>
   clone <profile>
   dot1-tlvs port-vlan|vlan-name
   dot3-tlvs link-aggregation|mac|mfs|power
   lldp-med-network-policy-profile <profile>
   lldp-med-tlvs capabilities|inventory|network-policy
```

```
   no ...
   optional-tlvs capabilities|management-address|port-description|system-description|system-
   name
   receive
   transmit
   transmit-hold <transmit-hold>
   transmit-interval <transmit-interval>
ap lldp med-network-policy-profile <profile>
   application-type guest-voice|guest-voice-signaling|softphone-voice|streaming-video|video-
   conferencing|video-signaling|voice|voice-signaling
   clone <profile>
   dscp <dscp>
   l2-priority <l2-priority>
   no ...
   tagged
   vlan <vlan>
```

The following commands create a LLDP MED network policy profile for streaming video applications and marks streaming video as high-priority traffic:

```
(host) (config) ap lldp med-network-policy-profile vid-stream
(host) (AP LLDP-MED Network Policy Profile "vid-stream") dscp 48
(host) (AP LLDP-MED Network Policy Profile "vid-stream")l2-priority 6
(host) (AP LLDP-MED Network Policy Profile "vid-stream")tagged
(host) (AP LLDP-MED Network Policy Profile "vid-stream")vlan 10
(host) (AP LLDP-MED Network Policy Profile "vid-stream")!
```

Next, the LLDP MED network policy profile is assigned to an LLDP profile, and the LLDP profile is associated with an AP wired-port profile:

```
(host) (config) ap lldp profile video1
(host) (AP LLDP Profile "video1")lldp-med-network-policy-profile vid-stream
(host) (AP LLDP Profile "video1")!
(host) (config)ap wired-port-profile corp2
(host) (AP wired port profile "corp2")lldp-profile video1
```

Advanced Voice Troubleshooting

ArubaOS enables you to debug voice issues more efficiently and quickly by providing detailed information about the voice calls, voice client status, and Call Detail Records (CDR). You can obtain the advanced troubleshooting information such as time of failure of the call, status of the client during the call failure, signal strength of the call, AP handoff information, and signaling message issues.

The following options allow you to easily troubleshoot voice call issues:

- View troubleshooting information on voice client status
- View troubleshooting information on voice call CDRs
- Debug voice logs
- View voice traces
- View voice configuration details

Viewing Troubleshooting Details on Voice Client Status

ArubaOS enables you to view the status of the voice clients. Additionally, it allows you to view more details such as AP handoff information and AP station report of an active call based on the client's IP address, or the MAC address.

The AP handoff information includes the AP events such as association request, re-association request, and de-authentication request with timestamps. The AP station report includes the AP MAC address, association time, average RSSI value, and retries count.

You can use the WebUI or CLI to view up to 60 entries of AP events and 30 entries of AP station reports for a voice client.

In the WebUI

1. Navigate to the **Monitoring > Voice > Voice Clients** page and select the voice client.
2. Click **HandOff Information** to view the AP station report and AP handoff information of the selected voice client.

In the CLI

To view the details of a voice client based on its IP address:

```
(host) #show voice client-status ip 10.15.20.63

Voice Client(s) Status
----------------------
Client(IP)    Client(MAC)         Client Name  ALG   Server(IP)    Registration State  Call
Status  BSSID              ESSID         AP Name   Flags
----------   -----------        -----------  ---   ----------    ------------------  ---------
--  -----              -----         -------   -----
10.15.20.63  00:00:f0:05:c9:e3  7812         h323  10.3.113.239  REGISTERED          In-Call
   00:0b:86:b7:83:91  st-voice-raj  RAP2-Lab  R
Num Clients:1
Flags: V - Visitor, W - Wired, R - Remote

AP Events
---------
Timestamp        BSS Id             Category  Event
---------        ------             --------  -----
Aug 13 09:22:57  00:0b:86:b7:83:91  Call      Call Start
Aug 13 11:29:34  00:0b:86:b7:83:91  Call      Call End
Aug 13 11:29:41  00:0b:86:b7:83:91  Call      Call Start
Aug 13 11:30:29  00:0b:86:b7:83:91  Call      Call End
Aug 13 11:30:39  00:0b:86:b7:83:91  Call      Call Start

AP Station Reports
------------------
Timestamp        BSS Id             RSSI  Tx      Tx-Drop  Tx-Data  Tx-Data-Retry  Tx-Data-
Bytes  Tx-Data-Time  Rx        Rx-Retry
---------        ------             ----  --      -------  -------  -------------  -----------
--  ------------  --        --------
Aug 13 12:35:05  00:0b:86:b7:83:91  61    253845  6904     253469   59805          22945603
   0             55171662  0

Current Active Calls
--------------------
Session Information                  Peer Party  Dir  Status     Dur(sec)  Orig time        R-
value  Codec  Band  Setup Time(sec)  Re-Assoc
-------------------                  ----------  ---  ------     --------  ---------        --
-----  -----  ----  ---------------  --------
10.15.20.56:3034 - 10.15.20.63:3140  -           IC   CONNECTED  3925      Aug 13 11:30:39  NA
      NA     NA    NA               0
```

To view the details of a voice client based on its MAC address:

```
(host) #show voice client-status sta 00:00:f0:05:c9:dc

Voice Client(s) Status
----------------------

Client(IP)    Client(MAC)         Client Name  ALG   Server(IP)    Registration State  Call
Status  BSSID              ESSID  AP Name  Flags
```

```
----------   -----------        -----------  ---   ----------    ------------------  ---------
--  -----               -----  -------  -----
10.15.20.56  00:00:f0:05:c9:dc  7811         sh323  10.3.113.239  REGISTERED          In-Call
    00:1a:1e:a8:2d:80  legap  W-AP65-2

Num Clients:1
Flags: V - Visitor, W - Wired, R - Remote


AP Events
---------
Timestamp        BSS Id             Category  Event
---------        ------             --------  -----
Aug 13 09:22:54  00:1a:1e:a8:2d:80  Call      Call Start
Aug 13 09:22:58  00:1a:1e:a8:2d:80  Call      Call End
Aug 13 09:26:22  00:1a:1e:a8:2d:80  Call      Call Start
Aug 13 11:29:33  00:1a:1e:a8:2d:80  Call      Call End
Aug 13 11:29:39  00:1a:1e:a8:2d:80  Call      Call Start
Aug 13 11:30:29  00:1a:1e:a8:2d:80  Call      Call End
Aug 13 11:30:36  00:1a:1e:a8:2d:80  Call      Call Start


AP Station Reports
------------------
Timestamp        BSS Id             RSSI  Tx      Tx-Drop  Tx-Data  Tx-Data-Retry  Tx-Data-
Bytes  Tx-Data-Time  Rx        Rx-Retry
---------        ------             ----  --      -------  -------  -------------  -----------
--  ------------  --        --------
Aug 13 12:38:03  00:1a:1e:a8:2d:80  44    795216  44158    794838   147824         78010395
   0             58366710  0

Current Active Calls
--------------------
Session Information                  Peer Party  Dir  Status     Dur(sec)  Orig time        R-
value  Codec  Band   Setup Time(sec)  Re-Assoc
-------------------                  ----------  ---  ------     --------  ---------        --
-----  -----  ----   ---------------  --------
10.15.20.63:3140 - 10.15.20.56:3034  -           OG   CONNECTED  4079      Aug 13 11:30:36  93
      NA     GREEN  NA               0
```

Viewing Troubleshooting Details on Voice Call CDRs

ArubaOS allows you to view the voice CDRs for the completed calls. Additionally, it enables you to view more details such as AP handoff information and AP station reports for a specific terminated call based on the CDR Id.

The AP handoff information includes the AP events such as association request, re-association request, and de-authentication request with timestamps. The AP station report includes the AP MAC address, association time, average RSSI value, and retries count.



ArubaOS pushes the generated CDRs to the syslog server to retain the older CDR data for a later analysis. The CDR data pushed to the syslog server do not contain the details of the AP stats and AP events.

You can use the WebUI or CLI to view the troubleshooting information on a voice call based on the CDR Id.

In the WebUI

1. Navigate to the **Monitoring > Voice > Call Detail Report** page.

This page displays the CDRs of the completed calls.

2. Click the **CDR Id** of a call to view the AP station reports, and the AP handoff information of the call.

In the CLI

To view the details of a completed call based on the CDR Id:

```
(host) #show voice call-cdrs cid 4

Voice Client(s) CDRs (Detail)
-----------------------------

CDR Id  Client IP    Client Name  ALG   Dir  Called/Calling Party  Status  Dur(sec)  Orig time
       R-value  Reason  Codec  Band    Setup Time(sec)  Re-Assoc  Initial-BSSID      Initial-
ESSID  Initial-AP Name
------  ---------    -----------  ---   ---  --------------------  ------  --------  ---------
       -------  ------  -----  ----    ---------------  --------  -------------      ---------
----  ---------------
4       10.15.20.62  3011         sccp  IC   3042                  SUCC    34        Aug 14
06:48:44  77               G711   YELLOW  0                1         00:1a:1e:a8:2d:80  legap
        W-AP65-2


AP Events
---------
Timestamp        BSS Id             Category       Event
---------        ------             --------       -----
Aug 14 06:48:53  00:1a:1e:a8:2d:80  AP Management  Assoc Req
Aug 14 06:48:53  00:1a:1e:a8:2d:80  AP Management  Assoc Resp


AP Station Reports
------------------
Timestamp        BSS Id             RSSI  Tx     Tx-Drop  Tx-Data  Tx-Data-Retry  Tx-Data-
Bytes  Tx-Data-Time  Rx     Rx-Retry
---------        ------             ----  --     -------  -------  -------------  ------------
-  ------------  --     --------
Aug 14 06:49:08  00:1a:1e:a8:2d:80  27    20466  6154     20460    2522           2310190
  0             26245  0
```

Enabling Voice Logs

ArubaOS allows you to debug voice logs. Additionally, it allows you to debug the voice logs for a specific voice client based on the client's MAC address.

You can use the WebUI or CLI to set the voice logging level to debugging.

In the WebUI

1. Navigate to the **Configuration > Management > Logging** page.
2. Click **Levels**.
3. Select the **voice** check box under the **User Logs** category.
4. Select **Debugging** from the **Log Level** drop down menu and click the **Done** button.

Figure 214 *Enable Voice Logging*

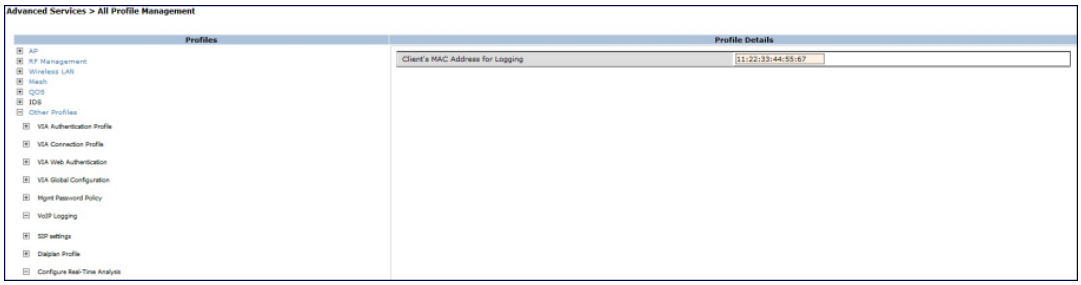


5. Click **Apply**.

Enabling Logging for a Specific Client

1. Navigate to the **Configuration > Advanced Services > All Profiles** page.
2. Expand **Other Profiles** under the **Profiles** section and click **VoIP Logging**.
3. Enter the MAC address of the voice client in the **Client's MAC address for logging** field.

Figure 215 *Enable Logging for a Voice Client*



4. Click **Apply** .

NOTE

To enable logging on a specific voice client, you must enable voice logs.

In the CLI

To set the voice logging level to debugging:

```
(host) #configure terminal
(config) #logging level debugging user subcat voice
```

To debug voice logs for a specific client:

```
(config) #voice logging
(VoIP Logging) #client-mac 11:22:33:44:55:67
```

To view the client's MAC address for logging:

```
(host) #show voice logging

VoIP Logging
------------
Parameter                          Value
---------                          -----
Client's MAC Address for Logging   11:22:33:44:55:67
```

Viewing Voice Traces

ArubaOS enables you to view the voice signaling message traces. You can view up to 8000 entries of trace messages. The trace message displays the ALG, client name, client's IP, event time, and the message direction. Additionally, it displays the BSSID information to help troubleshooting roaming issues.

You can use the WebUI or CLI to view the trace messages.

In the WebUI

1. Navigate to the **Monitoring > Voice > Voice Clients** page and select the voice client.
2. Click **Troubleshooting** to view the voice traces.

In the CLI

To view the voice signaling message traces:

```
(host)#show voice trace sip count 5
```

```
SIP Voice Client(s) Message Trace
---------------------------------
ALG  Client Name  Client(MAC)        Client(IP)    Event Time       Direction         Msg
           BSSID
---  -----------  -----------        ----------    ----------       ---------         ---
           -----
SIP  6202         00:03:2a:02:75:cc  10.15.20.123  Aug 14 13:14:32  Server-To-Client  200_OK
           00:0b:86:b7:83:91
SIP  6202         00:03:2a:02:75:cc  10.15.20.123  Aug 14 13:14:32  Client-To-Server  REGISTER
           00:0b:86:b7:83:91
SIP  6202         00:03:2a:02:75:cc  10.15.20.123  Aug 14 13:14:31  Server-To-Client  200_OK
           00:0b:86:b7:83:91
SIP  6202         00:03:2a:02:75:cc  10.15.20.123  Aug 14 13:14:31  Client-To-Server  REGISTER
           00:0b:86:b7:83:91
SIP  6202         00:03:2a:02:75:cc  10.15.20.123  Aug 14 13:14:29  Server-To-Client  4XX_
REQUEST_FAILURE  00:0b:86:b7:83:91
Num of Rows:5
```

Viewing Voice Configurations

ArubaOS allows you to view the details of the voice related configurations on your controller such as firewall policies, AP group profiles, SSID profiles, virtual AP group profiles, VoIP Call Admission Control profiles, 802.11k profiles, and SIP settings. Additionally, you can view the status of RTCP analysis, and SIP mid-call request timeout.



This release of ArubaOS does not support viewing the voice configuration details using the WebUI.

In the CLI

To view the voice configuration details on your controller:

```
(host) #show voice configurations

Voice firewall policies
-----------------------
Policy                      Action
------                      ------
Stateful SIP Processing     Enabled
Broadcast-filter ARP        Disabled

SSID Profiles
-------------
Profile Name            WMM       WMM-UAPSD  TSPEC Min Inactivity(msec)  ...  EDCA STA prof
EDCA AP prof  Strict SVP
------------            ---       ---------  --------------------------  ...  -------------  -
-----------  ----------
default                 Enabled   Enabled    100000                      ...  default
default       Disabled
qa-ma-vocera            Enabled   Enabled    0                                default
default       Disabled

AP Group Profiles
-----------------
Profile Name   VoIP CAC Profile
------------   ----------------
default        default
local          default

Virtual AP Group Profiles
-------------------------
```

```
Profile Name           802.11K Profile  HA Discovery on-assoc.  Drop Broadcast/Multicast
Broadcast ARP to Unicast
------------           ---------------  ----------------------  ------------------------  ----
------------------
abcd                   default          Disabled                Disabled
Disabled

VoIP Call Admission Control Profiles
------------------------------------
Profile Name  VoIP CAC
------------  ---------
default       Disabled

802.11K Profiles
----------------
Profile Name  Advertise 802.11K Capability
------------  ----------------------------
default       Disabled

SIP settings
------------
Parameter         Value
---------         -----
Session Timer     Disabled
Session Expiry    300 sec
Dialplan Profile  N/A

Voice rtcp-inactivity:disable
Voice sip-midcall-req-timeout:disable
```

AirGroup is a unique enterprise-class capability that leverages zero configuration networking to allow mobile device technologies, such as the AirPrint wireless printer service and the AirPlay mirroring service, to communicate over a complex access network topology.



The W-600 Series controllers do not support AirGroup.

Zero Configuration Networking

Zero configuration networking is a technology that enables service discovery, address assignment, and name resolution for desktop computers, mobile devices, and network services. It is designed for flat, single-subnet IP networks such as the home network of a user.

The suite of protocols introduced by Apple® for zero configuration networking over IP is referred to as Bonjour®. Bonjour is supported by most of the Apple product lines including the Mac OS X® operating system, iPhone®, iPod®, iPad®, Apple TV® and AirPort Express®. Bonjour is also included within popular software programs such as Apple iTunes®, Safari, and iPhoto®. Bonjour® can be installed on computers running Microsoft Windows® and is supported by most new network-capable printers.

Bonjour, the zero configuration implementation introduced by Apple, is supported by many Apple product lines, including devices using the OS X operating system, iPhone, iPod Touch®, iPad®, Apple TV, and AirPort Express. Bonjour is included with popular software programs such as Apple iTunes®, Safari, and iPhoto®. Bonjour can also be installed on computers running Microsoft Windows® and is supported by most new network-capable printers.

Bonjour locates devices such as printers, other computers, and the services offered by these devices by using multicast Domain Name System (mDNS) service records. Bonjour uses the link-scope multicast addresses, so each query or advertisement is limited to a specific VLAN. In large universities and enterprise networks, Bonjour capable devices connect to the network using different VLANs. As a result, an iPad on one enterprise VLAN will not be able to discover the Apple TV that resides on another VLAN. Broadcast and multicast traffic is filtered out of a wireless LAN network in an effort to reduce network traffic. This inhibits Bonjour (mDNS) services, which rely on multicast traffic.

ArubaOS supports DLNA (Digital Living Network Alliance); a network standard that is derived from UPnP (Universal Plug and Play) in addition to the mDNS protocol. DLNA uses the Simple Service Discovery Protocol (SSDP) for service discovery on the network. DLNA provides the ability to share digital media between multimedia devices, like Windows and Android, similar to how mDNS supports Zero Configuration Networking to Apple devices and services. ArubaOS ensures that DLNA seamlessly works with the current mDNS implementation. All the features and policies that are applicable to mDNS are extended to DLNA. This ensures full interoperability between compliant devices.

AirGroup Solution

AirGroup leverages key elements of Dell's solution portfolio including the ArubaOS software for Dell mobility controllers and Dell ClearPass Policy Manager (CPPM).

AirGroup performs the following functions:

- Enables users to discover network services across IP subnet boundaries in enterprise wireless and wired networks.
- Enables users to access the available AirGroup services such as AirPrint and AirPlay.
- Permits users to access conference room Apple TV during presentations, based on group-based access privileges.
- Provides and maintains seamless connectivity of clients and services across VLANs and SSIDs. It minimizes the mDNS traffic across the wired and wireless network, thereby preserving wired network bandwidth and WLAN airtime.

With AirGroup:

- An AirGroup operator—an end user such as a student can register personal devices. The devices registered by the operator can then automatically be shared with each other.
- Each user can create a user group, such as friends and roommates with whom the user can share the registered devices.
- AirGroup administrators can register and manage an organization's shared devices such as printers or conference room Apple TV. The administrator can grant global access to each device, or limit access based on user name, role, or location.

This chapter provides configuration information for network administrators to enable AirGroup on a Dell mobility controller and CPPM and to register devices with ClearPass Guest.

AirGroup also enables context awareness for services across the network:

- AirGroup is aware of personal devices. An Apple TV in a dorm room, for example, can be associated with the student who owns it.
- AirGroup is aware of shared resources, such as an Apple TV in a meeting room, a printer available to multiple users, or AirPlay in a classroom where a laptop screen is projected on HDTV monitor.
- AirGroup is aware of the location of services—for example, an iPad is presented with the closest printer location instead of all the printers in the building. If a user in a conference room wants to use an Apple TV receiver to project a MacBook screen on an HDTV monitor, the location-aware mobility controller shows the Apple TV that is closest to that user.

AirGroup Services

The AirGroup supports zero configuration services. The services are pre configured and are available as part of the factory default configuration. The administrator can also enable or disable individual services by using the controller WebUI.

The following services are enabled by default on the controller:

- AirPlay — Apple AirPlay allows wireless streaming of music, video, and slide shows from your iOS device to Apple TV and other devices that support the AirPlay feature.
- AirPrint — Apple AirPrint allows you to print from an iPad, iPhone, or iPod Touch directly to any AirPrint compatible printers.
- DIAL — Wi-Fi-enabled streaming devices like Google Chromecast, Roku, Amazon FireTV, and more advertise the Discovery and Launch (DIAL) protocol for clients to search for an available device on a wireless network. Once a device is discovered, the protocol synchronizes information on how to connect to the device. The streaming device connects to a television through an HDMI port to wirelessly stream video and music content to the TV screen from smart phone (both Android and Apple iOS), tablet, laptop or desktop computer devices.

The following services are disabled by default on the controller:

- iTunes — iTunes service is used by iTunes Wi-Fi sync and iTunes home-sharing applications across all Apple devices. For best practices, see the Apple iTunes Wi-Fi Synchronization and File Sharing on page 950.
- RemoteMgmt — Use this service for remote login, remote management, and FTP utilities on Apple devices.
- Sharing — Applications such as disk sharing and file sharing, use the service ID that are part of this service on one or more Apple devices. For best practices, see the Apple iTunes Wi-Fi Synchronization and File Sharing on page 950.
- Chat — The iChat (Instant Messenger) application on Apple devices uses this service.
- GoogleCast — Google Chromecast uses this service to stream video and music content from a smart phone to a TV screen using a wireless network. If this service is manually configured before the controller is upgraded to ArubaOS 6.4.1, the service continues to remain in the existing state.
- DLNA Media — Applications such as Windows Media Player use this service to browse and play media content on a remote device.
- DLNA Print — This service is used by printers which support DLNA.



AirGroup also supports custom and allowall services. For more information, see Integrated Deployment Model on page 954 and Integrated Deployment Model on page 954.

AirGroup Solution Components

AirGroup leverages key elements of Dell's solution portfolio that includes the ArubaOS software for Dell mobility controllers, CPPM, and ClearPass Guest. Table 203 describes the supported versions for each portfolio.

Table 203: *AirGroup Solution Component Supported Version*

Component	Minimum Version
ArubaOS (Mobility Controller)	6.4
CPPM and ClearPass Guest	6.0.2



It is recommended to use CPPM and ClearPass Guest version 6.3.

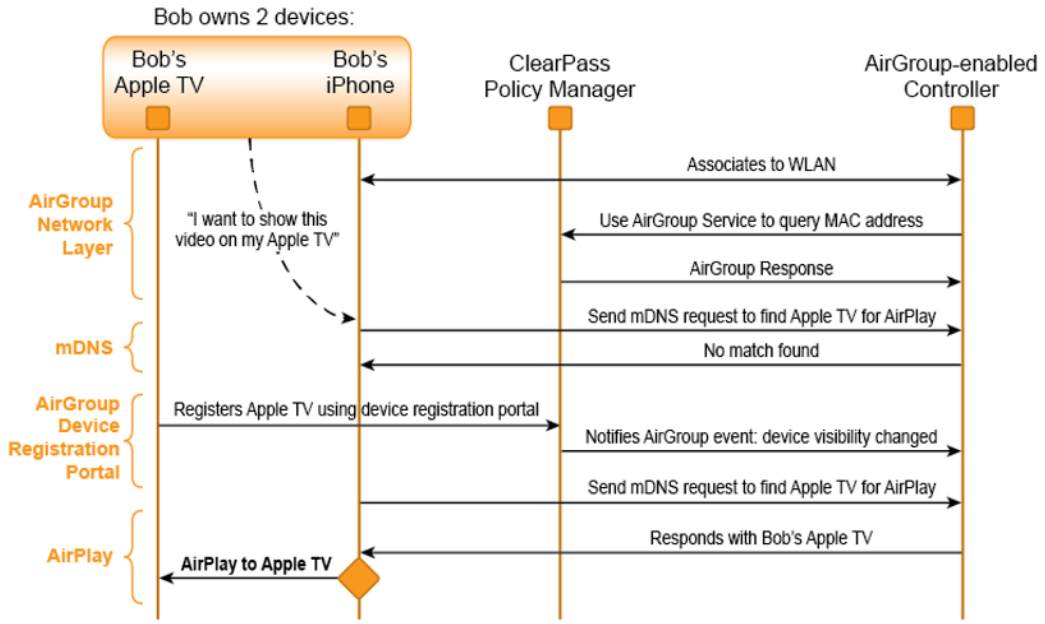
AirGroup and ClearPass Policy Manager

The AirGroup feature and CPPM work together to allow users to share personal devices.

- An AirGroup administrator uses ClearPass Policy Manager to authorize end users to register their personal devices.
- An AirGroup operator, an end user, registers devices (such as an Apple TV).
- Dell mobility controllers query ClearPass Policy Manager to associate the access privileges of each mobile device to its allowed services.

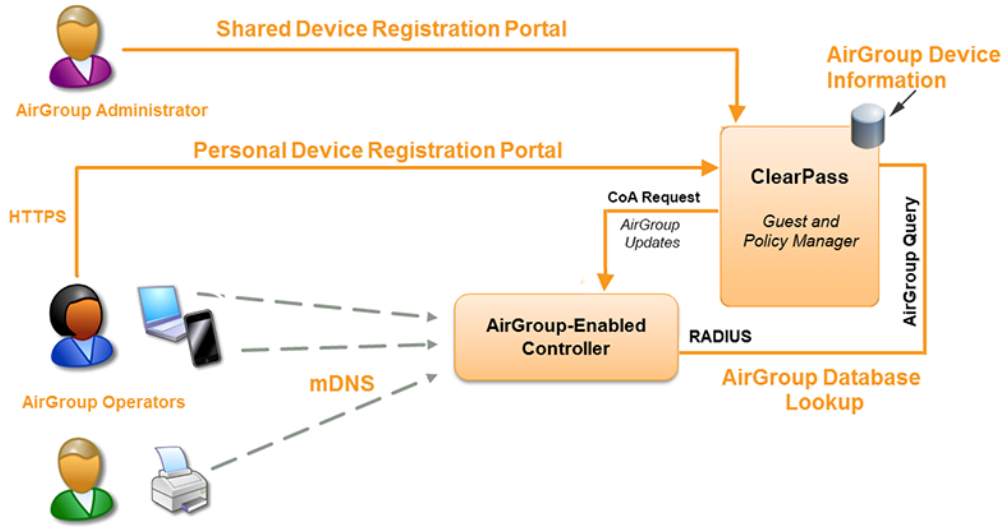
Figure 216 shows the AirGroup workflow that allows a user to register personal devices and use AirPlay to send an image from an iPhone to an Apple TV.

Figure 216 *AirGroup Enables Personal Device Sharing*



AirGroup enables context awareness for services across the network and supports a typical customer environment with shared, local, and personal services available to mobile devices. For example, in Figure 217, an AirGroup administrator registers the shared devices in ClearPass, and AirGroup operators register their personal devices in the ClearPass Guest portal. The AirGroup-enabled controller sends AirGroup queries to ClearPass for the registered devices' information. ClearPass sends the Change of Authorization (CoA) to notify the controller about the registered devices.

Figure 217 *AirGroup in a Typical Wireless Deployment*



AirGroup deployments that include both CPPM and an AirGroup controller support features that are described in AirGroup Services on page 941.

AirGroup Deployment Models

Integrated Deployment Model

In the integrated deployment model, AirGroup features are integrated with WLAN controllers that terminate APs and provide WLAN services. This deployment model also supports optional integration with ClearPass Policy Manager. If AirGroup is deployed in an integrated environment, you should upgrade all the controllers in your network to ArubaOS 6.4. For more information, see Integrated Deployment Model on page 954.

ArubaOS 6.4 supports a multi-controller AirGroup cluster. An AirGroup cluster consists of multiple controllers in various possible configuration combinations such as master-master, master-local, and local-local. If you are deploying AirGroup in a master-local topology with multiple local controllers that share the same user VLANs, use AirGroup in an integrated mode. Figure 218 shows an example of a master-local topology with shared, local, and personal services that are available to mobile devices. With AirGroup, the context-based policies determine the services visible to the end-user devices.

Figure 218 *Integrated AirGroup Network Topology*

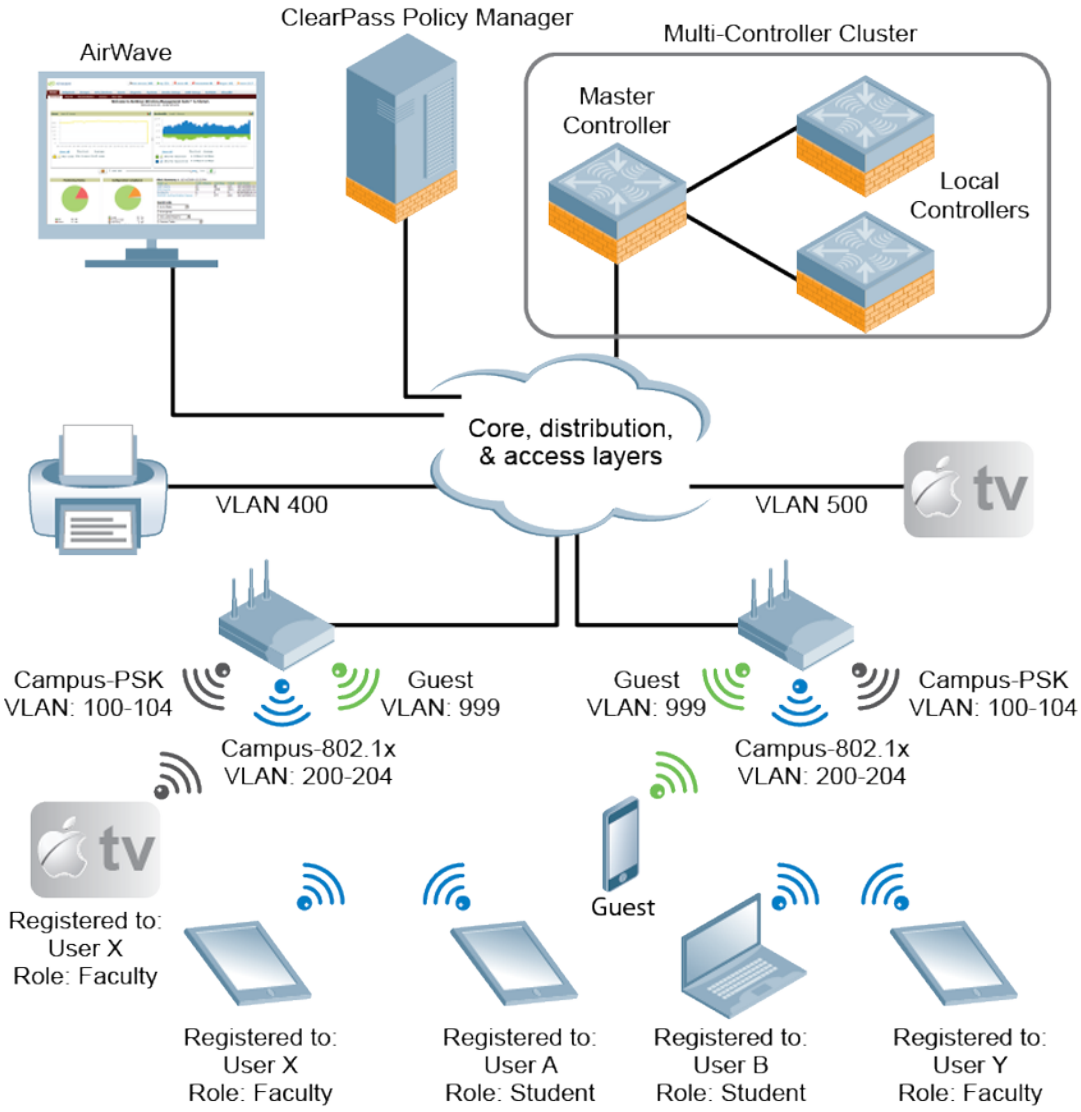


Table 204: *Sample policies for AirGroup*

mDNS Services	Faculty	Student	Visitor
	User X's iPad	**User B's MacBook**	**Windows Laptop**
Apple TV in the lab, registered to user role "Faculty"	Yes	No	No
Apple TV in the dorm room, registered to User B	No	Yes	No
Apple TV in a lecture hall accessible to Faculty	Yes	No	No
Printer located in a lab accessible to faculty and students	Yes	Yes	No

AirGroup with ClearPass Policy Manager

CPPM delivers identity and device-based network access control across any wired, wireless, and VPN infrastructure. AirGroup can be deployed with Dell ClearPass Policy Manager (recommended for large WLANs), or without ClearPass in smaller networks. If your deployment does not include ClearPass Policy Manager, features described in AirGroup Services on page 941 are not available.

Features Supported in AirGroup

The following AirGroup features are supported in ArubaOS:

Multi-Controller AirGroup Cluster

ArubaOS supports multiple mobility controllers running AirGroup to form a cluster. This feature enables iPad users on one controller to discover Apple TV available on another controller, if both controllers are part of the same cluster.

Multi-Controller AirGroup Cluster—Terminologies

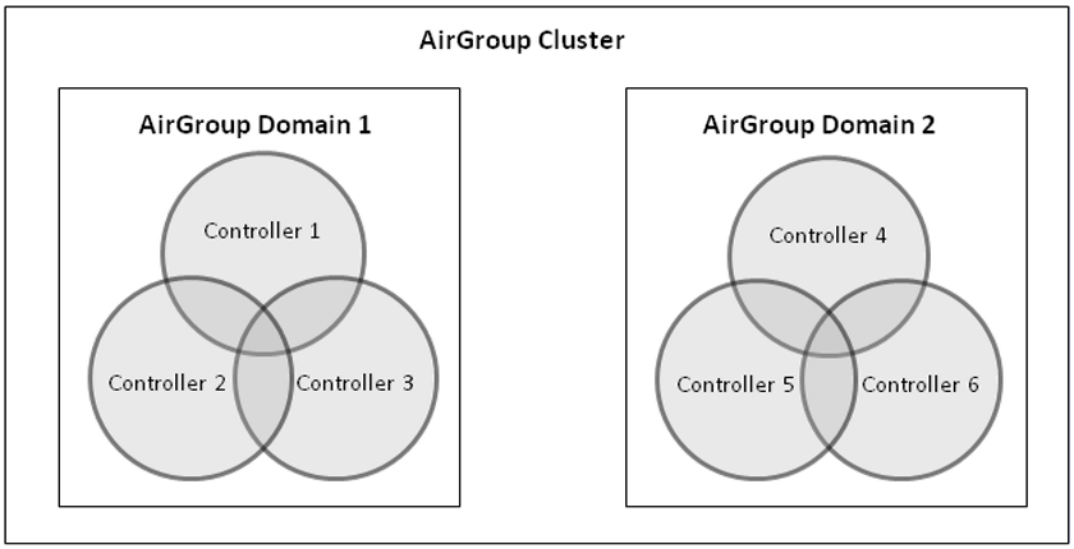
AirGroup Domain

An AirGroup domain is a set of controllers that are part of an AirGroup cluster. An administrator can configure multiple AirGroup domains for a site-wide deployment. Individual local controllers can independently select relevant multiple AirGroup domains to form a multi-controller AirGroup cluster.

AirGroup Cluster

One or several AirGroup domains can form an AirGroup cluster. AirGroup cluster can have 100 AirGroup domains. An AirGroup domain can include a list of likely controllers which may participate in the multi-controller AirGroup cluster. Figure 219 shows the AirGroup cluster and domain relationship:

Figure 219 *AirGroup cluster and domain relationship*



Active-Domain

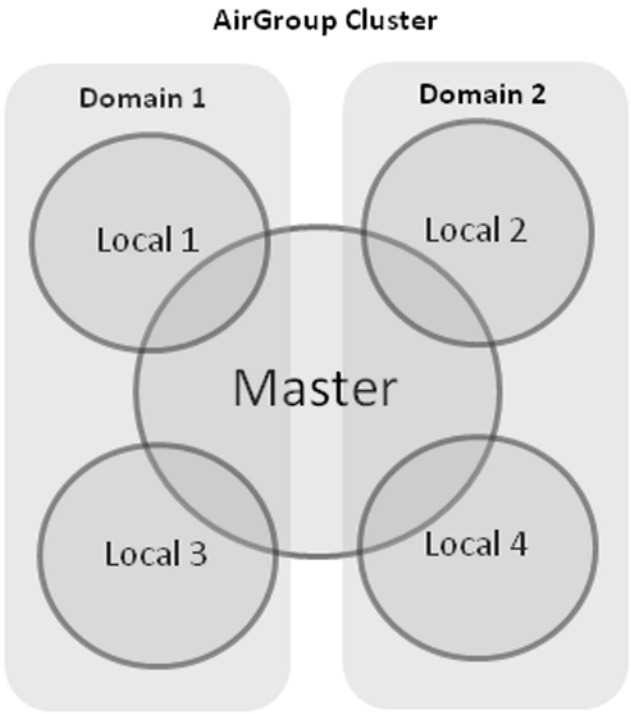
AirGroup allows one or more AirGroup domains to be a part of the AirGroup active-domain list on a controller. A master or local controller may participate in one or more AirGroup clusters based on its active-domain list. The mobility controller must set the corresponding domain as active for the controller to be part of the AirGroup cluster.

In Figure 219, Controller 1, 2, and 3 belong to **AirGroup Domain 1**.

Sample AirGroup Cluster Topology

Figure 220 shows a typical master-local multi-controller deployment. In this topology, four local controllers terminate on a single master controller.

Figure 220 *Typical Master-Local Multi-Controller Deployment*



Depends on the need, the administrator can configure the following topology:

Domain Definition

The administrator can define two domains with the following controllers in each domain:

- **Domain 1**: Local 1 (L1), Master (M), Local 3 (L3)
- **Domain 2**: Local 2 (L2), M, Local 4 (L4)

To configure an AirGroup domain, see Configuring an AirGroup Domain on page 960.

Active-Domain Definition

Based on the domain definition, each controller belongs to the following active-domain lists:

- **Active-Domain 1**: L1, M, L3
- **Active-Domain 2**: L2, M, L4

To configure an active domain, see Configuring an AirGroup active-domain on page 961.

AirGroup Controller Communication

Based on the domain and active-domain definitions, the AirGroup controller communication takes place in the following manner:

- L1, M, and L3 can communicate with each other as they are part of active-domain 1.
- L2, M, and L4 can communicate with each other as they are part of active-domain 2.
- M can communicate with L1, L2, L3, and L4 as M is part of active-domain 1 and 2.
- L1 and L3 cannot communicate with L3 and L4, because they do not have a common active-domain and they do not share the same VLAN.

AirGroup Server Discovery

- iPad users in L1, M, and L3 can discover any Apple TV or AirPrint Printer in L1, M, and L3.

- iPad users in L2, M, and L4 can discover any Apple TV or AirPrint Printer in L2, M, and L4.
- iPad users in M can discover any Apple TV or AirPrint printer in L1, L2, L3, and L4 and vice-versa.
- iPad users in L1 and L3 cannot discover any Apple TV or AirPrint printer in L2 and L4 and vice-versa.

Scalability

In a multi-controller deployment, there is a scaling limit of 2,000 AirGroup servers and 16,000 AirGroup users for all controllers in a cluster. If you require more servers and users than the prescribed limit, configure multiple clusters so that each cluster is within the prescribed limit. For detailed scalability information, see AirGroup Scalability Limits on page 952.

An AirGroup domain can include a list of controllers, which may participate in the multi-controller AirGroup cluster. Depending on the deployment setup, the IP address in the AirGroup domain could either be the controller IP or VRRP IP address. The configuration elements are defined by the administrator on a master controller and its associated local controllers that share the same configuration. The actual AirGroup multi-controller cluster may include one or several local controllers, and this cluster is defined by including one or several relevant AirGroup domains, on the respective local controller, in the active-domain list. As a result, a master or local controller may participate in one or more AirGroup clusters based on its active-domain list.

Incorrect or incomplete configuration of the controllers participating in an AirGroup cluster can lead to disjoint clusters. In a disjoint cluster, an AirGroup user will not have a seamless view of the AirGroup servers spanning multiple controllers. Therefore ensure that the participating controllers in an AirGroup cluster are configured appropriately.

The AirGroup domain configurations are restricted to the master controller. This ensures all local controllers in a master-local setup have unique AirGroup domain names. If duplicate AirGroup domain names on multiple master controllers are encountered, ensure that the duplicate AirGroup domain names have the same values to participate in a single AirGroup cluster.



Any controller that shares VLANs with another controller must be part of the same AirGroup multi-controller cluster.

When an AirGroup controller has the list of all the controllers in the multi-controller table, it uses a Dell proprietary protocol called Process Application Programming Interface (PAPI) to communicate with other controllers in the table. The PAPI control channel carries AirGroup specific packets only. For configuration details, see Configuring an AirGroup Domain on page 960.

Master-Local Controller Synchronization

Administrators can configure AirGroup from the master controller to ease deployment. The master controller then synchronizes the AirGroup configuration elements with all the local controllers it manages. For more information, see Master-Local Controller Synchronization on page 954.

Pre-configured AirGroup Services

The following services are pre-configured and available as part of the factory default configuration:

- AirPlay
- AirPrint
- iTunes
- RemoteMgmt
- Sharing
- Chat
- GoogleCast

- DIAL
- DLNA Print
- DLNA Media



DIAL is enabled by default. DLNA Print and DLNA Media are disabled by default.

For more information, see Integrated Deployment Model on page 954.

AirGroup IPv6 Support

A mobility controller supports IPv6 enabled users (for example, iPad) and servers (Apple TV, AirPrint printers). All the AirGroup features are available for both IPv4 and IPv6 clients. On any dual stack client, you must restart the client if the IPv4 interface is disabled.

Limitations

IPv6 support is limited to AirGroup users and servers only. The IPv4 addresses are supported only in the following scenarios:

- When forming an AirGroup cluster, only IPv4 controller addresses are supported.
- AirGroup supports IPv4 RADIUS clients only.



The controller can identify any IPv6 AirGroup servers, only when they proactively advertise their services.

To enable or disable AirGroup IPv6 support on the controller, see Enabling or Disabling AirGroup Global Setting on page 955.

DLNA UPnP Support

AirGroup supports DLNA (Digital Living Network Alliance); a network standard that is derived from UPnP (Universal Plug and Play) in addition to the mDNS protocol. For more information, see Zero Configuration Networking on page 940.

AirGroup mDNS Static Records

AirGroup provides the ability for an administrator to create mDNS static records as group and individual records and add them to cache. For more information, see AirGroup mDNS Static Records on page 973.

Group Based Device Sharing

AirGroup supports **User Group** and this is an add-on to the existing device sharing mechanisms such as user-name, user-role, and location based device sharing using CPPM. For more information, see Group Based Device Sharing on page 971.

Dashboard Monitoring Enhancements

- The AirGroup service names in the **AirGroup** row are clickable in the **AirGroup** section of the **Dashboard > Usage** page of the WebUI. If you click the service name, you are redirected to the **Dashboard** > **AirGroup** page which displays a list of AirGroup servers filtered by Service Name.
- In the **Dashboard > Clients** page, the **AirGroup** column is added to display the devices that are listed as mDNS, DLNA or both. If a device does not support both mDNS and DLNA, this field is blank.
- The following enhancements are added in the **Dashboard** > **AirGroup** page of the WebUI:

- A new **AirGroup type** column is added and this column specifies if the type of the AirGroup device is mDNS, DLNA or both.
- The MAC address of each AirGroup user and server is now clickable. If you click MAC link, you are redirected to the **Dashboard** > **Clients** > **Summary** page > **AirGroup** tab. If an AirGroup user or AirGroup server is a wired trusted client, the MAC address is not clickable.

ClearPass Policy Manager and ClearPass Guest Features

With CPPM portal for WLAN administrators, you can register shared device such as conference room Apple TV and printers. The ClearPass Guest portal for WLAN users allows end users to register their personal devices. For more information on AirGroup configuration on CPPM, see the *ClearPass Policy Manager User Guide* and *ClearPass Guest Deployment Guide*.

Best Practices and Limitations

Apple iTunes Wi-Fi Synchronization and File Sharing

When the controller receives mDNS response for a service, the controller caches such records and does not propagate to other users. But for services like iTunes Wi-Fi synchronization and File Sharing to work seamlessly, such mDNS responses must be propagated to other users on the controller even if they do not query for it.

To ensure that applications such as iTunes Wi-Fi synchronization and File Sharing work seamlessly, ArubaOS selectively forwards these mDNS responses to AirGroup users, based on the user-name CPPM policy of the AirGroup server. Hence, for a customer to use these services, it is necessary to configure user-name based CPPM policies for the AirGroup devices.

Firewall Configuration

The following firewall configuration settings are recommended:

Disable Inter-User Firewall Settings

Some firewall settings can prevent the untrusted clients from communicating with each other. When these settings are enabled, an untrusted client such as an iPad may not be able to send its image to an Apple TV on the same controller.

Use the following commands to disable the virtual AP global firewall options and allow Bonjour services to use AirGroup.

- **no firewall deny-inter-user-bridging**
- **no firewall deny-inter-user-traffic**
- **no ipv6 firewall deny-inter-user-bridging**

ValidUser ACL Configuration

The **ValidUser** Access Control list (ACL) must allow mDNS packets with the source IP as a link local address. Do not use a **ValidUser** ACL if the user VLAN interfaces of the AirGroup controller are not configured with an IP address.

Allow GRE and UDP 5353

mDNS discovery uses the predefined port UDP 5353. If there is a firewall between the AirGroup controller and WLAN controller, ensure that your firewall policies allow GRE and UDP 5353. DLNA uses the predefined port UDP 1900.

Recommended Ports

The ArubaOS role-based access controls for wireless clients use ACLs to allow or deny user traffic on specific ports. Even though mDNS discovery uses the predefined port UDP 5353, application-specific traffic for services like AirPlay may use dynamically selected port numbers. As a best practice, add or modify ACLs to allow traffic on the ports as described in Table 205 and Table 206.



AirPlay operates using dynamic ports, however, printing protocols like AirPrint use fixed ports.

Ports for AirPlay Service

Enable the following ports for the AirPlay services.

Table 205: *Ports for AirPlay Service*

Protocol	Ports
TCP	• 5000 • 7000 • 7100 • 8612 • 49152-65535
UDP	• 7010 • 7011 • 8612 • 49152-65535

Ports for AirPrint Service

Enable the following ports to allow AirGroup devices to access AirPrint services.

Table 206: *Ports for AirPrint Service*

Protocol	Print Service	Port
TCP	Datastream	9100
TCP	IPP	631
TCP	HTTP	80
TCP	Scanner	9500
TCP	HTTP-ALT	8080

AirGroup Services for Large Deployments

All Bonjour services are enabled in AirGroup by default. Large deployments with many wireless and wired users often support a large number of advertised Bonjour services, which can consume a significant amount of system resources. For large scale deployments, enable the **AirPlay** and **AirPrint** services, disable the **allowall** service, and then block all other Bonjour services. See Integrated Deployment Model on page 954 for a complete list of AirGroup configuration options.

AirGroup Scalability Limits

Table 207 displays the total number of AirGroup servers (Apple TV, AirPrint printer) and users (iPad) supported in individual controllers:

Table 207: *AirGroup Server and User Limits in Controller*

Controller Model	Number of AirGroup servers	Number of AirGroup users
W-3200	500	1500
W-3400	1000	3000
W-3600	2000	6000
W-6000M3	2000	6000
W-7005	300	700
W-7010	500	1500
W-7030	1000	3000
W-7210	2000	9000
W-7220	2000	12000
W-7240	2000	16000



In a multi-controller deployment, there is a scaling limit of 2,000 AirGroup servers and 16,000 AirGroup users for all controllers in a cluster. If you require more servers and users than the prescribed limit, configure multiple clusters, so that each cluster is within the prescribed limit.

The ArubaOS scaling limits are based on the following metrics:

- Memory Utilization
- CPU Utilization

Memory Utilization

The memory utilization is affected by the number of AirGroup servers and users in an AirGroup cluster. In an AirGroup cluster, the total number of AirGroup servers and users cannot exceed the limit defined by the top-end controller. For example, in an AirGroup cluster of one W-3200 controller and two W-6000M3 controllers, the cluster limit is determined as per the scaling limit of the top-end controller which is the W-6000M3 controller. For the W-3200 controller in the cluster, the controller platform limit of the W-3200 controller is applied. Based on the memory utilization, Table 207 summarizes the maximum number of AirGroup servers and users for all supported controller platforms.

CPU Utilization

The CPU utilization is measured by the rate at which the controller receives mDNS packets. The rate of mDNS packets in the cluster depends on the number of AirGroup servers, users, and number of applications installed on these devices. The rate of mDNS packets handled by the supported controller platform varies. Table 208 displays the total number of mDNS packets received per second by supported the controller platforms:

Table 208: *mDNS Packet Limits in Controller*

Controller Model	mDNS packets per second (pps)
W-3200	10
W-3400	10
W-3600	20
W-6000M3	20
W-7005	20
W-7010	40
W-7030	50
W-7210	20
W-7220	25
W-7240	30

Use the following command to determine the number of mDNS packets received per second by the controller:

```
show airgroup internal-state statistics
```



Issue this command multiple times to measure the time difference and the mDNS packet count.

General AirGroup Limitations

The AirGroup feature has the following limitations:

- AirGroup is supported only in tunnel and decrypt-tunnel forwarding modes.
- If you use CPPM to define AirGroup users, the shared user and role lists, and location attributes cannot exceed 1000 characters.
- The RTSP protocol does not support AirPlay on an Apple TV receiver if you enable NAT on the user VLAN interface.
- The location-based access feature only supports AP FQLNs (Fully Qualified Location Names) configured in the format **<ap name>.floor <number>.<building>.<campus>**. The AP names cannot contain periods.
- AirGroup's DLNA discovery works across VLANs, however, media streaming from Windows Media Server does not work across VLANs. This limitation is because of Digital Rights Management (DRM) support in Windows Media Server, which restricts media sharing across VLANs. Media streaming works only when both client and server are connected to the same VLAN.
- Android devices cannot discover media server while using the native music and video player applications and when they are connected across VLANs. For example, Samsung Tab 3 cannot discover the media server on Samsung Galaxy S4 while using the native music and video player applications. Android devices can discover media server when they are connected in the same VLAN. This restriction is forced by Samsung devices.
- Xbox cannot be added as an extender to the Windows clients using the Windows Media Center application with the AirGroup feature enabled. You need to disable the AirGroup feature to add Xbox as an extender.

- Wireless Clients such as iPad and iPhone running the Sonos Controller application cannot discover Sonos music system with the AirGroup is enabled.

Integrated Deployment Model

In the integrated deployment model, AirGroup features are integrated with the WLAN controller that terminates all APs and provides WLAN services. This deployment model also supports optional integration with CPPM. When you implement AirGroup in an integrated deployment, upgrade the controller to ArubaOS 6.4 or later, and trunk all VLANs with wired devices (such as printers) to the AirGroup controller.



If your deployment requires ClearPass Policy Manager integration, complete the procedures described in *ClearPass Policy Manager User Guide* and *ClearPass Guest Deployment Guide* before performing the steps described in this section.

Master-Local Controller Synchronization

You can configure AirGroup from the master controller to ease the deployment. The master controller then synchronizes the AirGroup configuration elements on all associated local controllers it manages. AirGroup configurations can belong to any of the following categories:

Master — These commands must be configured from a master controller. The master controller pushes the AirGroup configurations to all the applicable local controllers.

- AirGroup custom service definition. For more information, see Integrated Deployment Model on page 954.
- AirGroup disallow user-role (service filtering) definition. For more information, see Configuring the disallow-role for an AirGroup Service on page 958.
- AirGroup disallow VLAN (service filtering) definition. For more information, see Restricting AirGroup Servers on a VLAN based on an AirGroup Service on page 959.
- AirGroup CPPM enforce registration. For more information, see Configuring the AirGroup-CPPM Interface on page 965.
- AirGroup controller-CPPM Interface definition. For more information, see Configuring the AirGroup-CPPM Interface on page 965.
- AirGroup multi-controller domain definition. For more information, see Configuring an AirGroup Domain on page 960.
- AirGroup CPPM query interval definition. For more information, see Configuring the CPPM Query Interval on page 965.

Local — There are a few configuration limitations on the local controller. The local controller can only include the existing AirGroup domains in the AirGroup active-domain list, applicable for this controller. The local controller cannot define or edit an AirGroup domain.

These configuration commands are applicable to both master and local controllers. The master controller does not push the following AirGroup configuration commands to all applicable local controllers.

- AirGroup enable/disable parameter. For more information, see Enabling or Disabling AirGroup Global Setting on page 955.
- AirGroup service enable/disable parameter. For more information, see Enabling or Disabling an AirGroup Service on page 959.
- AirGroup allowall service status. For more information, see Integrated Deployment Model on page 954.
- AirGroup disallow VLAN (global) definition. For more information, see Restricting AirGroup Servers for a VLAN on page 958.
- AirGroup multi-controller active-domain definition. For more information, see Configuring an AirGroup active-domain on page 961.

Configuring an AirGroup Integrated Deployment Model

Use the following procedures to enable the AirGroup feature and configure AirGroup services.

Enabling or Disabling AirGroup Global Setting

Starting from ArubaOS 6.4, AirGroup is disabled by default. To configure AirGroup global parameters, use the following procedure:

Using the WebUI

To enable or disable the AirGroup global setting using the controller WebUI:

1. Navigate to **Configuration > Advanced Services > AirGroup**.
2. Select the **AirGroup Settings** tab.
3. Under **Global Setting > AirGroup Status**, select **enable** from the drop-down list to enable the AirGroup feature.
4. Under **Global Settings** > **AirGroup MDNS Status**, select **enabled** from the drop-down list to enable the MDNS.
5. Under **Global Settings** > **AirGroup DLNA Status**, select **enabled** from the drop-down list to enable the DLNA.
6. Under **Global Setting > AirGroup CPPM enforce registration**, select **enable** from the drop-down list to register an AirGroup server on a CPPM server.

   For more information on AirGroup CPPM enforce registration, see Configuring the AirGroup-CPPM Interface on page 965.
7. Under **Global Setting > AirGroup IPV6 Support**, select **enable** from the drop-down list.



The global AirGroup status must be enabled on the controller to enable AirGroup IPv6 support. For more information, see AirGroup IPv6 Support on page 949.

8. Under **Global Setting > AirGroup CPPM query interval**, enter a value in the range of 1 to 24 hours.

   The default value is 10. For more information on AirGroup CPPM query interval, see Configuring the CPPM Query Interval on page 965.
9. Under **Global Setting > AirGroup location discovery**, select **enable** from the drop-down list.

   If enabled, AirGroup user can discover shared devices based on the user's proximity to the AirGroup server. If disabled, location based filtering does not apply. Users can discover far servers. For more information on location attributes in CPPM, see Table 209.
10. Under **Global Setting > AirGroup Active Wireless Discovery**, select **enable** from the drop-down list.

    If enabled, AirGroup controller actively sends refresh requests to discover wireless servers. If disabled, the controller sends refresh requests to wired AirGroup servers only.
11. Click **Apply**.



AirGroup CPPM enforce registration, **AirGroup CPPM query interval**, **AirGroup location discovery**, and **AirGroup Active Wireless Discovery** parameters are available on the master controller only. The master controller pushes these configurations to all the applicable local controllers.

Table 209 shows the location attributes a device can register with CPPM and the corresponding behavior on the controller:

Table 209: *Location Attributes in CPPM*

Location Attribute	Tag=Value Format	Description
AP-Name based	ap-name=<name>	When the location is set to **ap-name**, all AirGroup users connected to this AP and other APs that are in the same RF neighborhood can access the shared device.
AP-Group based	ap-group=<group>	When the location attribute is set to **ap-group**, all AirGroup users associated to the APs in the specified AP group can access the shared device.
AP-FQLN based	fqln=<fqln>	When the location attribute is set to **ap-FQLN**, all AirGroup users connected to APs on the same floor, and to the APs on a floor above or below the configured APs can access the shared device.

Using the CLI

Access the controller's command-line interface and use the following command to enable or disable the AirGroup Global Setting:

```
(host) (config) #airgroup {enable | disable}
(host) (config) #airgroup cppm-server enforce-registration
(host) (config) #airgroup ipv6
(host) (config) #airgroup query-interval <1..24>
(host) (config) #airgroup location-discovery {enable | disable}
(host) (config) #airgroup active-wireless-discovery {enable | disable}
```

Enabling or Disabling mDNS and DLNA

You can enable and disable mDNS and DLNA using CLI commands and WebUI.

Using the CLI

Use the following command to enable or disable the mDNS or DLNA for an AirGroup service:

```
airgroup [mdns|dlna] enable|disable
```



Both mDNS and DLNA are disabled by default.

Use the following command to view the status of mDNS and DLNA features:

```
#show airgroup status
```

Viewing AirGroup Global Setting on Controller

Using the WebUI

To view the global setting of AirGroup in the controller WebUI:

1. Navigate to **Configuration > Advanced Services > AirGroup**.
2. Select the **AirGroup Settings** tab to view the AirGroup **Global Setting** in the controller.

Using the CLI

Use the following command to view the global settings of the AirGroup configuration and AirGroup services configured in your WLAN controller.

```
(host) #show airgroup status
```

For more information, see *Dell Networking W-Series ArubaOS 6.4 Command-Line Interface Reference Guide*.

Defining an AirGroup Service

The AirGroup solution defines the concept of configurable AirGroup services. One or more mDNS and DLNA services can be configured on the mobility controller. When you define an mDNS service as an AirGroup service, you can implement policies to restrict its availability to a specific user role or VLAN.

The following services are preconfigured and available as part of the factory default configuration:

- AirPlay
- AirPrint
- iTunes
- RemoteMgmt
- Sharing
- Chat
- GoogleCast
- DIAL
- DLNA Media
- DLNA Print

Using the WebUI

An administrator can configure and use up to 100 AirGroup services, and each AirGroup service can support up to 100 service elements. To define an AirGroup service using the controller WebUI:

1. Navigate to **Configuration > Advanced Services > AirGroup**.
2. On the **AirGroup service details** tab, click **Add New**.
3. Enter the name of the AirGroup profile in the **Name** field.
4. Enter the description for the AirGroup profile in the **Description** field.
5. Select **Enable** to enable this service.
6. Enter the VLANs that need to be restricted in the **Disallow VLANs** field.
7. Enter the roles that need to be restricted in the **Disallow Roles** field.
8. Enter the Service ID of the AirGroup service in the **Services IDs** field.
9. Click **OK** and then click **Apply**.

Table 210 describes the configuration parameters of an AirGroup service:

Table 210: *AirGroup Service Parameters*

Parameter	Description
Name	Name of the AirGroup Service.
Description	Enter the description for the AirGroup Service.
Enable	Enables the AirGroup service.

Parameter	Description
Disallow VLANs	User VLANs restricted from accessing the service.
Disallow Roles	User Roles restricted from accessing the service.

Service IDs	Specifies the mDNS or DLNA service IDs.
	An AirGroup mDNS service ID is the name of a Bonjour service offered by a Bonjour-enabled device or application. Bonjour defines mDNS service ID strings using the following format **<underscore>servicename<period><underscore>protocol.local**
	Example: `_airplay._tcp.local`
	The mDNS service ID string is case sensitive and must be entered as is without any modification, with the exception of the .local portion of the service ID which is optional.
	When you add an existing mDNS service ID to a new service, Airgroup automatically deletes the mDNS service ID from the old service and displays a warning message. A sample warning message is as follows:
	`service id <_ssh._tcp> removed from <remotemgmt> and added to <remotelogin>`
	The DLNA service IDs are colon separated and the service ID should have the following format to discover DLNA server or devices with the maximum label size of 128 characters:
	`urn:domain-name:device:deviceType:ver`
	`urn:domain-name:service:serviceType:ver`
	For example, you can use the following service ID to support DLNA media server under AirGroup:
	`urn:schemas-upnp-org:device:MediaServer:1`
	NOTE: Cache refresh mechanism is not required for DLNA, as the DLNA devices advertise their service periodically.

Using the CLI

Use the **airgroupservice** command to define an AirGroup service using the command-line interface.

```
airgroupservice <name>
```

Sample Configuration

The following example configures the **iPhoto** service with access to the **_dpap._tcp** service ID to share photos across MacBooks:

```
(host) (config) #airgroupservice iPhoto
(host) (config-airgroupservice) #description "Share Photos"
(host) (config-airgroupservice) #id _dpap._tcp
```

Configuring the disallow-role for an AirGroup Service

An AirGroup service is accessible to all user devices associated to your controller by default. The **disallow-role** parameter prevents devices with the specified user roles from accessing AirGroup services.

```
airgroupservice <string>
   disallow-role <string>
```

Sample Configuration

```
(host) (config) #airgroupservice iPhoto
(host) (config-airgroupservice) #disallow-role guest
```

Restricting AirGroup Servers for a VLAN

An AirGroup service is accessible to user devices in all VLANs configured on your controller by default. Use the following command to enable or disable AirGroup access to devices in a specific VLAN:

```
airgroup vlan <VLAN ID> {allow | disallow}
```

Sample Configuration

```
(host) (config) #airgroup vlan 5 disallow
```

Restricting AirGroup Servers on a VLAN based on an AirGroup Service

To prevent user devices on a specific VLAN from accessing a specific AirGroup service, use the disallow-vlan option.

```
airgroupservice <string>
  disallow-vlan <string>
```

Sample Configuration

```
(host) (config) #airgroupservice airplay
(host) (config-airgroupservice) #disallow-vlan 5
```

Viewing AirGroup Disallowed VLAN Policy Details

Use the following command to view the status of a disallowed VLAN policy.

```
show airgroupservice [dlna|mdns] [verbose]
```

Viewing An AirGroup Disallowed VLAN

Use the following command to view the status of the disallowed AirGroup VLANs:

```
show airgroup vlan
```

For more information, see *Dell Networking W-Series ArubaOS 6.4 Command-Line Interface Reference Guide*.

Enabling the allowall Service

The **allowall** service is a preconfigured AirGroup service that enables the controller to permit all AirGroup services by default, without requiring an administrator to configure an AirGroup service.

Using the WebUI

Use the following steps to enable the **allowall** service using the controller WebUI:

1. Navigate to **Configuration > Advanced Services > AirGroup**.
2. In the **AirGroup service details** tab, select the checkbox next to the **allowall** service and click **Enable**. To disable this service, select the **allowall** checkbox and click **Disable**.
3. Click **Apply**.

Using the CLI

Use the following command to enable or disable the allowall service:

```
airgroup service allowall {enable | disable}
```

Sample Configuration

```
(host) (config) #airgroup service allowall enable
```

Enabling or Disabling an AirGroup Service

Using the WebUI

To enable or disable an AirGroup service using the controller WebUI:

1. Navigate to **Configuration > Advanced Services > AirGroup**.
2. On the **AirGroup service details** tab, select the AirGroup service and click **Enable** or **Disable**.
3. Click **Apply**.

Using the CLI

Use the following command to enable or disable an AirGroup service:

```
airgroup service <string> {enable | disable}
```

Sample Configuration

```
(host) (config)#airgroup service airplay disable
```

Viewing AirGroup Service Status

Using the WebUI

Use the following steps to view the status of AirGroup services using the controller WebUI:

1. Navigate to **Configuration > Advanced Services > AirGroup**.
2. Under the **AirGroup service details** tab, view the status of all the AirGroup services.

Using the CLI

Use the following command to verify the status of an AirGroup Service:

```
show airgroup status
```

Sample Configuration

For sample configuration, see show airgroup status.

Viewing Blocked Services

The **airgroup service <servicename> disable** command blocks an AirGroup service by blocking the service IDs for that service. When you enable an AirGroup service, service IDs of that service are enabled automatically. To view the list of blocked services, use the **show airgroup blocked-service-id** command.

Using the CLI

```
show airgroup blocked-service-id [mdns|dlna]
```

Viewing AirGroup Service Details

Using the WebUI

To view the AirGroup service details using the controller WebUI:

1. Navigate to **Configuration > Advanced Services > AirGroup service details**.
2. Under **AirGroup service details** tab, click on any of the service name to view the service details.

Using the CLI

Use the following command to view the service details of all AirGroup services:

```
show airgroupservice
```

Sample Configuration

For sample configuration, see Viewing AirGroup Disallowed VLAN Policy Details on page 959.

Configuring an AirGroup Domain

An administrator can configure multiple AirGroup domains for a site-wide deployment. Individual local controller can independently choose relevant multiple AirGroup domains to form a multi-controller AirGroup cluster.



An administrator can configure and use up to 100 AirGroup domains, and each AirGroup domain can support up to 100 IP addresses.

NOTE

A domain can be configured only on a master controller only. However, active domains can be added/removed on any controllers.

The following procedure configures a cluster of mobility controllers for a domain:

Using the WebUI

1. Navigate to **Configuration > Advanced Services > AirGroup**.
2. Select the **AirGroup Settings** tab.
3. Under the **AirGroup Domains** section, click **Add New**.
4. In the **Name** field, enter the domain name.
5. In **Description** field, enter a short description of the domain name.
6. Select the **Active** checkbox to enlist the domain in the active-domain list of a controller.
7. Under the **IP Address** section, enter the controller or VRRP IP to be a part of this domain and click **Add**.

NOTE

If the deployment includes master or local redundancies, use the VRRP IP address in the domain definition. Else, use the controller IP address.

8. Click **Ok** and **Apply**.

Using the CLI

```
[no] airgroup domain <string>
   [no] ip-address <A.B.C.D>
   [no] description <string>
```

Sample Configuration

```
(host) (config) #airgroup domain Campus1
(host) (config-airgroup-domain) #ip-address 10.10.10.1
(host) (config-airgroup-domain) #ip-address 11.11.11.1
(host) (config-airgroup-domain) #description AirGroup_campus1
```

Viewing an AirGroup Domain

The following procedure displays a list of AirGroup domains configured:

Using the WebUI

1. Navigate to **Configuration > Advanced Services > AirGroup**.
2. Select the **AirGroup Settings** tab. The list of AirGroup domains are displayed under the **AirGroup Domains** section.

Using the CLI

```
show airgroup domain
```

Configuring an AirGroup active-domain

AirGroup allows one or more AirGroup domains to be a part of the AirGroup active-domain list of a controller. A master or local controller may participate in one or more AirGroup cluster based on its active-domain list. The mobility controller must set the corresponding domain as active for the controller to be part of the AirGroup cluster.

The following procedure configures an AirGroup active-domain for AirGroup cluster:

Using the WebUI

For the WebUI procedure, see Configuring an AirGroup Domain on page 960.

Using the CLI

```
[no] airgroup active-domain <string>
```

Sample Configuration

```
(host) (config) #airgroup active-domain campus1
(host) (config) #airgroup active-domain campus2
```

Viewing an AirGroup active-domains

The following procedure displays a list of AirGroup active-domains configured:

Using the WebUI

1. Navigate to **Configuration > Advanced Services > AirGroup**.
2. Select the **AirGroup Settings** tab. The **Active-Domain** and **Status** column displays a list of AirGroup active-domains under the **AirGroup Domains** section.

Using the CLI

```
show airgroup active-domains
```

Viewing AirGroup VLAN Table

The following procedure displays the disallowed AirGroup VLANs.

Using the WebUI:

1. Navigate to **Configuration > Advanced Services > AirGroup**.
2. Select the **AirGroup Settings** tab. The list of disallowed AirGroup VLANs are displayed under the **VLAN Table** section.

Using the CLI

For the CLI command, see Viewing An AirGroup Disallowed VLAN on page 959

Viewing AirGroup Multi-Controller Table

All controllers communicate with each other based on the multi-controller table in an AirGroup cluster. This table is a combination of controllers specified in each domain, as part of active-domains.

The following command displays the IP address of all the controllers participating in an AirGroup multi-controller environment:

Using the CLI

```
show airgroup multi-controller-table
```

Controller Dashboard Monitoring

The **Dashboard > Usage** page of the WebUI has an additional **AirGroup** section, which displays all the AirGroup services available and number of servers offering the service. It is aggregated by the total number of AirGroup servers sorted by the services they advertise.

Figure 221 *AirGroup Dashboard Usage*

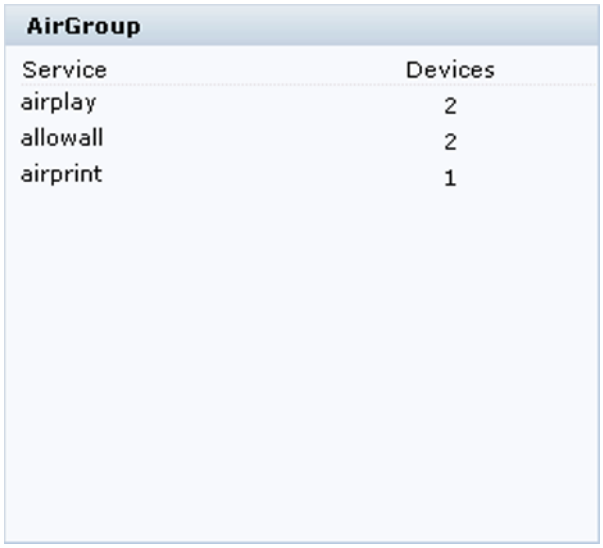


Table 211: *AirGroup Dashboard Usage*

Column	Description
Service	Displays the services advertised by AirGroup servers discovered by the controller.
Devices	Displays the number of AirGroup servers advertising a particular service.

Click the **IP** link to view the client details in the **Dashboard > Clients** page of the WebUI. The client details page has a tab called **AirGroup**. The **AirGroup** tab in the details page displays a list of all the far and near end devices that are either accessible or not accessible by the specific client.



In a multi-controller topology, only AirGroup clients and servers that are connected to the same controller are listed under the near or far devices categories. AirGroup does not fetch this information from other controllers that are part of the same multi-controller domain.

- A device is classified as a **Near Device** if it is registered with CPPM and the location is set to any one of the following:
  - AP-Group same as the client
  - AP-FQLN same as the client
  - AP-FQLN corresponding to adjacent floors of the client
  - AP-Name same as the client
  - AP-Name which is an RF neighbor of the client

  For more information on location, see Location Attributes in CPPM.
- A device is classified as a **Far Device** if none of the above criteria is met. Devices that are neither registered nor have a location defined in CPPM are classified as **Far Devices** by default.
- A device is classified as **Accessible** or **Non Accessible** based on the CPPM policies and disallow-role configuration.

Figure 222 *Near and Far Accessible Devices*



Table 212: *Near and Far Accessible Devices*

Column	Description
MAC Address	Displays the MAC address of the near or far AirGroup server that is accessible by an AirGroup client.
Name	Displays the hostname of the near and far AirGroup server that is accessible by an AirGroup client.
Service	Displays the AirGroup service advertised by an AirGroup server.

Figure 223 *Near and Far Non Accessible Devices*

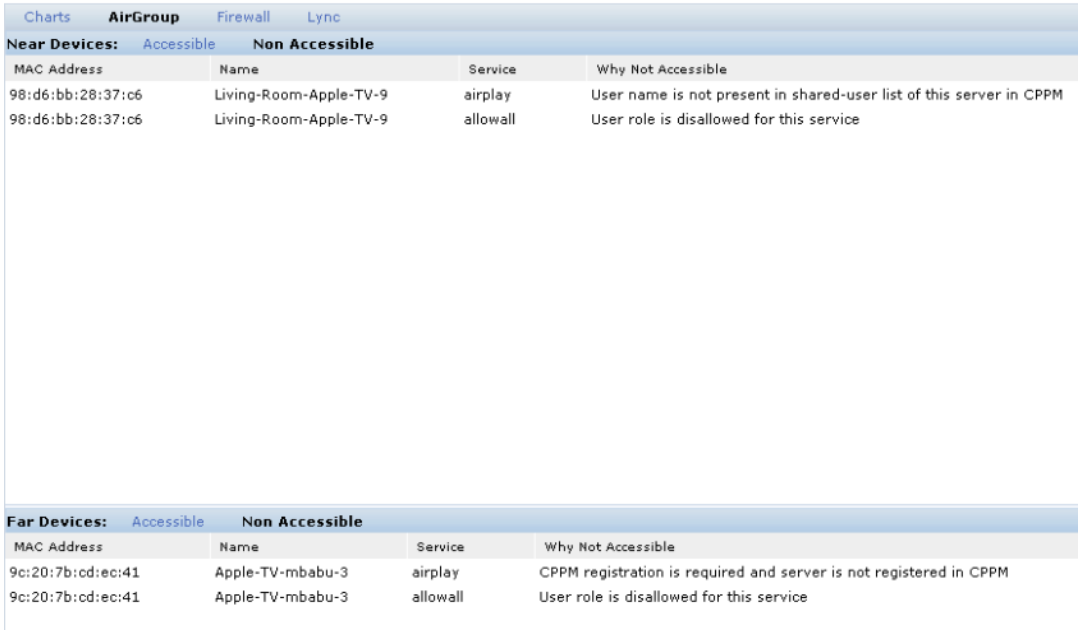


Table 213: *Near and Far Non Accessible Devices*

Column	Description
MAC Address	Displays the MAC address of the near or far AirGroup server that is not accessible by an AirGroup client.
Name	Displays the hostname of the near and far AirGroup server that is not accessible by an AirGroup client.
Service	Displays the AirGroup service advertised by an AirGroup server.
Why Not Accessible	Displays the reason for not accessing the AirGroup server.

Configuring the AirGroup-CPPM Interface

Configure the AirGroup and ClearPass Policy Manager (CPPM) interface to allow an AirGroup controller and CPPM to exchange information about the owner, visibility, and status for each mobile device on the network. The following procedures configure the AirGroup-CPPM interface:

- Configuring the CPPM Query Interval on page 965
- Defining a CPPM and RFC3576 Server on page 966
- Assigning CPPM and RFC 3576 Servers to AirGroup on page 969
- Viewing the CPPM Server Configuration on page 970
- Configuring the AirGroup-CPPM Interface on page 965

Configuring the CPPM Query Interval

The AirGroup CPPM query interval refreshes the CPPM entries at periodic intervals. The minimum value is 1 hour and the maximum value is 24 hours. The default value is 10 hours.

Using the WebUI

1. Navigate to the **Configuration > Advanced Services > AirGroup** page.

2. Select the **AirGroup Settings** tab.
3. Under **Global Setting > AirGroup CPPM query interval**, enter a value in the range of 1 to 24 hours.
4. Click **Apply**.

Using the CLI

```
[no] airgroup cppm-server query-interval <1..24>
```

Sample Configuration

```
(host) (config) #airgroup cppm-server query-interval 9
```

Viewing the CPPM Query Interval

The following procedure displays the configured CPPM query interval value.

Using the WebUI

1. Navigate to the **Configuration > Advanced Services > AirGroup** page.
2. Select the **AirGroup Settings** tab. The **AirGroup CPPM query interval** displays the value in hours under the **Global Setting** section.

Using the CLI

```
show airgroup cppm-server query-interval
```

Sample Configuration

```
(host) #show airgroup cppm-server query-interval

CPPM Server Query Interval
--------------------------
Timer Value  Unit
-----------  ----
9            hours
```

The output of this command includes the following information:

Table 214: *show airgroup cppm-server query-interval*

Column	Description
`Timer Value`	Displays the number of hours.
`Unit`	Displays the unit in hours.

Defining a CPPM and RFC3576 Server

You must define one or more CPPM servers to be used by the AirGroup RADIUS client, and an RFC 3576 (dynamic authorization) server. If multiple CPPM servers are defined, the servers are listed in a sequential order. The AirGroup RADIUS client will use the first available server on this list.

Table 215 describes the configuration parameters for a CPPM server.

Table 215: *CPPM Server Configuration Parameters*

Parameter	Description
Host	IP address or fully qualified domain name (FQDN) of the authentication server. The maximum supported FQDN length is 63 characters.
Key	Shared secret between the controller and the authentication server. The maximum length is 128 characters.
Authentication Ports	Authentication port on the server. Default: 1812
Accounting Ports	Accounting port on the server. Default: 1813
Retransmits	Maximum number of retries sent to the server by the controller before the server is marked as down. Default: 3
Timeout	Maximum time, in seconds, that the controller waits before timing out the request and resending it. Default: 5 seconds
NAS ID	Network Access Server (NAS) identifier to use in RADIUS packets.
NAS IP	NAS IP address to send in RADIUS packets. You can configure a "global" NAS IP address that the controller can use for communications with all CPPM servers. However, that the controller will only use this global NAS IP if you do not configure a server-specific NAS IP. To set the global NAS IP in the WebUI, navigate to the **Configuration > Security > Authentication > Advanced** page. To set the global NAS IP in the CLI, use the **ip radius nas-ip <A.B.C.D>** command.
Source Interface	Enter a VLAN number ID. This value allows you to use source IP addresses to differentiate RADIUS requests, and associates a VLAN interface with the RADIUS server to allow the server-specific source interface to override the global configuration. • If you associate a Source Interface (by entering a VLAN number) with a configured server, the source IP address of the packet will be the same as the IP address of the interface. • If you do not associate the Source Interface with a configured server (leave the field blank), the IP address of the global Source Interface will be used.
Use MD5	Use a MD5 hash of the cleartext password.
Use IP address for calling station ID	Select this checkbox to use an IP address instead of a MAC address for the calling station ID.
Mode	Enables or disables the server.

Configuring a CPPM Server

You can configure a CPPM server for AirGroup using the WebUI or CLI.



Server-derived user roles or VLANs configured in this server group are not applicable to AirGroup.

Using the WebUI

To configure a CPPM server using the controller WebUI:

1. Navigate to **Configuration > Security > Authentication > Servers**.
2. Select **Radius Server** to display the CPPM Server List.
3. To configure a CPPM server, enter the name for the server and click **Add**.
4. Select the name to configure server parameters. Select the **Mode** checkbox to activate the authentication server.
5. Click **Apply**.

Using the CLI

Use the following commands to configure a CPPM server using the CLI:

```
aaa authentication-server radius <name>
   host <ipaddr>
   key <key>
   enable
```

Sample Configuration

```
(host)(config) #aaa authentication-server radius emp_accounts
(host)(RADIUS Server "emp_accounts") #host 10.100.8.32
(host)(RADIUS Server "emp_accounts") #key employee123
(host)(RADIUS Server "emp_accounts") #enable
```

Configuring the CPPM Server Group

Using the WebUI

To configure a CPPM server group using the controller WebUI:

1. Navigate to **Configuration > Security > Authentication > Servers**.
2. Select **Server Group** to display the Server Group list.
3. Enter the name of the new server group and click **Add**.
4. Select the name to configure the server group.
5. Under **Servers**, click **New** to add a server to the group.
   a. Select a server from the drop-down list and click **Add Server**.
   b. Repeat the above step to add other servers to the group.
6. Click **Apply**.

Using the CLI

Use the following commands to configure a CPPM server group using the CLI:

```
aaa server-group <name>
auth-server <name>
```

Sample Configuration

```
(host) (config) #aaa server-group employee
(host) (Server Group "employee") #auth-server emp_accounts
```

Configuring an RFC 3576 Server

Using the WebUI

To configure an RFC 3576 server by using the controller WebUI:

1. Navigate to **Configuration > Security > Authentication > Servers**.
2. Select **RFC 3576 Server**.

3. Enter the **IP address** and click **Add**.
4. Select the IP address to enter the shared secret key in the **Key** text box.
5. Retype the shared secret key in the **Retype** text box.

Using the CLI

Use the following commands to configure an RFC 3576 server using the CLI:

```
aaa rfc-3576-server <server_ip>
   key <string>
```

Sample Configuration

```
(host) (config) #aaa rfc-3576-server 10.100.8.32
(host) (RFC 3576 Server "10.100.8.32") #key employee123
```

Assigning CPPM and RFC 3576 Servers to AirGroup

Use the following procedures to assign CPPM and RFC 3576 servers to AirGroup.



An AirGroup RFC 3576 server cannot use the same port as an authentication module RFC 3576 server. To avoid conflicts, use a non-standard port for the AirGroup RFC 3576 server.

Using the WebUI

Use the following procedure to configure the AirGroup AAA profile by using the controller WebUI:

1. Navigate to **Configuration > Advanced Services > All Profiles**.
2. Expand the **Other Profiles** menu and select **AirGroup AAA Profile**.
3. In the **Configure dead time for a down Server** text box in the **Profile Details** window, enter a maximum period in minutes, so that a client that does not send user traffic for the given period is considered idle.
4. Enter the UDP port number in the **Configure UDP port to receive RFC 3576 server requests** field. If your network uses an RFC 3576 server for authentication, select a different port for the AirGroup 3576 server. The default in ClearPass Guest is 5999.



In this release of ArubaOS, the user-defined UDP port number for RFC3576 server is automatically permitted by the firewall. The administrator does not have to explicitly define a firewall policy to permit this port.

5. Identify the AirGroup CPPM server group. In the **Profiles** list, select the **Server Group** under the **AirGroup AAA Profile** menu.
6. In the **Profile Details** window, click the **Server Group** drop-down list to select the desired CPPM server group.
7. Click **Apply**.
8. Identify the RFC 3576 server. In the **Profiles** list, select **RFC 3576 Server** under the **AirGroup AAA Profile** menu.
9. Enter the IP address of the RFC 3576 server in the **Add a profile** text box.
10. Click **Add** and **Apply**.

Using the CLI

Execute the following commands to configure the AirGroup AAA profile using the CLI:

```
airgroup cppm-server aaa
   rfc-3576-server <ip address>
   rfc-3576_udp_port <port number>
   server-dead-time <time>
```

```
server-group <server group name>
```



If your network uses an RFC 3576 server for authentication, select a different port for the AirGroup 3576 server. The default port in ClearPass Guest is 5999.

Sample Configuration

```
(host) (config) # airgroup cppm-server aaa
(host) (Airgroup AAA profile) #rfc-3576-server 10.15.16.25
(host) (Airgroup AAA profile) #rfc3576_udp_port 21334
(host) (Airgroup AAA profile) #server-dead-time 10
(host) (Airgroup AAA profile) #server-group employee
```

Viewing the CPPM Server Configuration

Using the WebUI

To view the CPPM server configuration by using the controller WebUI:

1. Navigate to **Configuration > Advanced Services > AirGroup**.
2. Under the **AirGroup Settings** tab, the **AirGroup CPPM server aaa** section displays the CPPM Server configuration.;

Using the CLI

Use the following CLI command to view data for the ClearPass Policy Manager servers:

```
(host) #show airgroup cppm-server aaa

Config-cppm-server-aaa
----------------------
Parameter                                                Value
---------                                                -----
Server Group                                             RADIUS_4
RFC 3576 server                                          Test1
Configure dead time for a down Server                    10
Configure UDP port to receive RFC 3576 server requests   N/A
```

The output of this command includes the following information:

Table 216: *show airgroup cppm-server aaa*

Column	Description
`Parameter`	Displays the AAA parameters for AirGroup.
`Value`	Displays the value entered for each AAA parameter.

Verifying CPPM Device Registration

Use the following command to display information for devices registered in ClearPass Policy Manager.

```
(host) #show airgroup cppm entries
```

For more information, see *Dell Networking W-Series ArubaOS 6.4 Command-Line Interface Reference Guide*.

Configuring CPPM to Enforce Registration

The AirGroup solution allows the users to view all mDNS devices by default. AirGroup provides a set of policy definitions to allow or disallow one of more AirGroup servers from being visible to specific AirGroup users.

If an AirGroup server is not registered on a CPPM server, by default, the server will be visible to all AirGroup users. The administrator must register an AirGroup server to allow or disallow this server from being visible to specific AirGroup users.

The following procedure registers an AirGroup server on a CPPM server:

Using the WebUI

To configure using the controller WebUI:

1. Navigate to **Configuration > Advanced Services > AirGroup**.
2. Select the **AirGroup Settings** tab.
3. Under **Global Setting > AirGroup CPPM enforce registration**, select **Enabled** from the drop-down list.
4. Click **Apply**.

Using the CLI

Use the following command to force AirGroup servers to register with CPPM. This option is disabled by default:

```
(host) (config) #airgroup cppm-server enforce-registration
```

To verify the CPPM Registration Enforcement status, use the following command:

```
(host) #show airgroup status
```

For more information, see *Dell Networking W-Series ArubaOS 6.4 Command-Line Interface Reference Guide*.

Group Based Device Sharing

Dell Networking W-Series ArubaOS 6.4.x AirGroup supports sharing AirGroup devices such as AppleTV, Printer, and so on to a **User Group** using CPPM. This is an add-on to the existing device sharing mechanisms such as username, user-role, and location based device sharing. A **User Group** is a logical association of users.

A user can be a part of groups that are defined in Active Directory. User group attribute for each user in a controller is learnt, when a user is associated to wireless network. In ArubaOS, this is initially learnt in auth module (authentication process). Auth module sends RADIUS request to RADIUS server as a part of 802.1x authentication and the RADIUS server fetches the user group attribute in the form of vendor specific attribute (VSA) from the Active Directory. Subsequently, AirGroup obtains this information from Auth module. This is similar to user's role, however, a user can be a part of more than one groups.

When AirGroup learns about a new device, it interacts with ClearPass Guest to obtain the shared attributes. Starting from Dell Networking W-Series ArubaOS 6.4.x, the shared group(s) attribute is also obtained along with the following attributes:

- Device owner
- Shared location(s)
- Shared user(s)
- Shared role(s)



The group based device sharing feature is supported in CPPM 6.3 and higher versions.



A user can be a part of maximum 32 user groups. This needs to be defined as comma separated string in Active directory. Each group name can contain a maximum of 63 characters and the entire group name strings cannot exceed 320 characters.

The AirGroup policy engine is enhanced to compare the user's group membership (obtained using auth module) and shared groups to determine if a user can discover the specific AirGroup server or not.

Example

The following example displays the status of the AirGroup server (Apple TV, AirPrint Printer, Google ChromeCast, and so on) in a controller:

```
(host) #show airgroup servers


AirGroup Servers
----------------
MAC                IP            Type  Host Name  Service     VLAN  Wired/Wireless
---                --            ----  ---------  -------     ----  --------------
5c:3c:27:14:6e:01  10.15.121.240  mDNS             airplay      2    wireless

Role           Group        Username  AP-Name
----           -----        --------  -------
authenticated  Mathematics  Mike      104_AP105

Num Servers: 1, Max Servers: 2000.
```

The following example displays the shared group information for devices registered in ClearPass Guest:

```
(host) #show airgroup cppm entries

ClearPass Guest Device Registration Information
-----------------------------------------------
Device             device-owner  shared location-id AP-name  shared location-id AP-FQLN
------             ------------  -------------------------  --------------------------
00:1e:65:2d:ae:44  N/A

shared location-id AP-group  shared user-list  shared group-list  shared role-list
---------------------------  ---------------  -----------------  ----------------
                                                 Physics

CPPM-Req  CPPM-Resp
--------  ---------
1         1

Num CPPM Entries:1
```

The following example describes the user Alice is the member of Mathematics group and hence cannot discover the 00:1e:65:2d:ae:44 appleTv (specified in the example above) as it is not shared with the Mathematics group. Similarly, the user Bob can view the appleTv as it is shared with the Physics group.

```
(host) #show airgroup users

AirGroup Users
--------------
MAC                IP            Type  Host Name  VLAN  Role           Group        Username
AP-Name
---                --            ----  ---------  ----  ----           -----        --------
-------
74:e1:b6:15:25:7e  10.15.121.240  mDNS  iPad-358   2     authenticated  Mathematics  Alice
104_AP105
```

```
b0:65:bd:09:b6:79  10.15.121.240  mDNS            2     authenticated  Physics      Bob
104_AP105

Num Users: 2, Max Users: 6000.
```

AirGroup mDNS Static Records

AirGroup processes mDNS packets advertised by servers and creates the relevant cache entries. When a query comes from a user, AirGroup responds with the appropriate cache entries with the relevant policies applied. Starting with ArubaOS 6.4, AirGroup provides the ability for an administrator to add the mDNS static records to cache, when a server is:

- not mDNS compliant.
- connected to a VLAN that is not trunked to the AirGroup supported controller.

The administrator can add these records manually to the cache using CLI commands for the servers that adhere to the above conditions.

Important Points to Remember

Remember the following points when you create mDNS static records and add them to the cache:

- The mDNS static records do not expire as there is no cache refresh for static records. These static records can be deleted by an administrator.
- The Administrator needs to ensure that the relevant records are updated manually, when the IP address of a server is changed.
- The **Disallow role** policy configured on the CLI is accepted for static records. The **Disable service** policy is accepted while responding to a query. Administrator has the privilege to configure static records of a disabled service. **Disallow vlan** is not applicable for static records.
- ClearPass Policy Manager policies work with static servers.

Creating mDNS Static Records on a Controller

The Administrator can create the static records using the following methods:

- Group mDNS static records
- Individual mDNS static records

Group mDNS Static Records

You can create a group of mDNS records for a device. This section describes how to create static records of a server as a group using the CLI.

Creating a PTR Record

Use the following command to create a PTR record:

```
(config) # airgroup static mdns-record ptr <mac_addr> <mdns_id> <domain_name> [server_ipaddr]
(config-airgroup-record) #
```



After creating a PTR record, controller displays the **(config-airgroup-record) #** prompt and you can create SRV, A, AAAA, and TXT records under this prompt.



After creating a PTR, SRV, TXT, A, and AAAA static record, you can use the **show airgroup cache entries** command to view and verify the records created. You can view only the static records in the output of the **show airgroup cache entries static** command.

Creating an SRV Record

Use the following command to create an SRV record:

```
(config-airgroup-record) # srv <port> <priority> <weight> <host_name>
```

Creating an A Record

Use the following command to create an A record:

```
(config-airgroup-record) #a <ipv4addr>
```



You can create/delete an A record if a corresponding SRV record is available.

Creating an AAAA Record

Use the following command to create an AAAA record:

```
(config-airgroup-record) #aaaa <ipv6addr>
```



You can create/delete an AAAA record if a corresponding SRV record is available.

Creating a TEXT Record

Use the following command to create a TEXT record:

```
(config-airgroup-record) #txt <text>
```

Individual Static mDNS Records

You can create individual static records independently for each record type.

Creating an Individual SRV Record

Use the following command to configure an individual SRV record:

```
airgroup static mdns-record srv <mac_addr> <domain_name> <port> <priority> <weight> <host_
name> [ server_ipaddr]
```

Creating an Individual TEXT Record

Use the following command to configure an individual TEXT record:

```
airgroup static mdns-record txt <mac_addr> <domain_name> <text> [server_ipaddr]
```

Creating an Individual A Record

Use the following command to configure an individual A record:

```
airgroup static mdns-record a <mac_addr> <host_name> <ipv4addr> [server_ipaddr]
```

Creating an Individual AAAA Record

Use the following command to configure an individual AAAA record:

```
airgroup static mdns-record aaaa <mac_addr> < host_name> <ipv6addr> [server_ipaddr]
```



You can delete the mDNS records by appending `no` at the beginning of the command. Ensure that the **[server_ipaddr]** parameter is not added while deleting mDNS records.

For more information, see *Dell Networking W-Series ArubaOS 6.4 Command-Line Interface Reference Guide*.

Troubleshooting and Log Messages

Controller Troubleshooting Steps

Use the following procedure to prevent potential errors in a controller:

1. Execute the **show airgroup internal-state statistics** CLI command and ensure that the **Sibyte Messages Sent/Recv** counters increment over a period of time.
2. Enable mDNS logs using the **logging level debugging system process mdns** command, and capture the output of **show log system all** when the issue occurs. Review any obvious error print statements.
3. Save the output of **show airgroup cache entries** and **show airgroup cppm entries** and look for any discrepancies.

ClearPass Guest Troubleshooting Steps

ClearPass Guest includes AirGroup-related events in the application log files. You can configure logging levels to provide debugging information.

To show debugging information in event logs:

1. In ClearPass Guest, go to **Administration > AirGroup Services** and click the **Configure AirGroup Services** command link. The **Configure AirGroup Services** form opens.
2. In the **AirGroup Logging** drop-down list, select either **Debug—log debug information** or **Trace—log all debug information**. When one of these options is selected, debugging information is provided in the events log.
3. Click **Save Configuration**.

For up-to-date information, see the *ClearPass Guest Deployment Guide*.

ClearPass Policy Manager Troubleshooting Steps

Monitoring and reporting services in ClearPass Policy Manager provide insight into system events and performance.

To show incoming AirGroup requests from the controller:

1. In ClearPass Policy Manager, navigate to **Monitoring > Live Monitoring > Access Tracker**. The **Access Tracker** list view opens.
2. Click an event's row to view details. The **Summary** tab of the **Request Details** view opens. Additional details may be viewed on the **Input**, **Output**, or **Alerts** tabs, or you can click **Show Logs** to view logging details.

For up-to-date information, see the *ClearPass Policy Manager User Guide*.

Log Messages

Display AirGroup logs by issuing the following commands in the controller CLI:

- **show log all**
- **show log system all**
- **show log user all**
- **show log user-debug all**

The log debug messages for the mDNS process are not enabled by default. To enable specific logging levels, use the following CLI commands in configuration mode:

To enable high level mDNS debug messages:

```
(host)(config) #logging level debugging system process mdns
```

To enable mDNS packet processing messages:

```
(host)(config) #logging level debugging system process mdns subcat messages
```

To enable mDNS CLI configuration messages:

```
(host)(config) #logging level debugging system process mdns subcat configuration
```

To enable mDNS Auth and CPPM user messages:

```
(host)(config) #logging level debugging user process mdns
```

Show Commands

Use the following show commands to view AirGroup configuration data and statistics in the controller:

Viewing AirGroup mDNS and DLNA Cache

```
show airgroup cache entries [mdns|dlna|static]
```

Viewing AirGroup mDNS and DLNA Statistics

```
show airgroup internal-state statistics [mdns|dlna]
```

Viewing AirGroup VLANs

```
(host) #show airgroup vlan
```

Viewing AirGroup Servers

Use the following command to view the AirGroup server (Apple TV, AirPrint Printer, Google ChromeCast, and so on) status in the controller:

```
show airgroup servers [dlna|mdns] [verbose]
```

Viewing AirGroup Users

```
show airgroup users [mdns|dlna] [verbose]
```

Viewing Service Queries Blocked by AirGroup

This command displays the service ID that was queried but not available in the AirGroup service table.

```
show airgroup blocked-queries [mdns|dlna]
```

Viewing Blocked Services

The **`airgroup service <servicename> disable`** command disables an AirGroup service by blocking the service IDs for that service. When you enable an AirGroup service, service IDs of that service are enabled automatically. To view the list of blocked services, use the following command:

```
show airgroup blocked-service-id [mdns|dlna]
```

AirGroup Global Tokens

In an AirGroup network, AirGroup devices generate excess mDNS query and response packets. Using airgroup global-credits command, the AirGroup controller restricts these packets by assigning tokens. The controller processes these mDNS packets based on this token value. The controller rejects any packets beyond this token limit. The token renews every 15 seconds. The renewal interval is not a configurable parameter.

In the following example, the AirGroup controller restricts the number of query packets to 450 and response packets to 90 from AirGroup devices in a time frame of 15 seconds.

```
(host)(config) #airgroup global-credits 450 90
```

The following command displays tokens assigned to query and response packets. It displays the current and user configured global tokens.

```
(host) #show airgroup global-credits
```

For more information, see *Dell Networking W-Series ArubaOS 6.4 Command-Line Interface Reference Guide*.

Chapter 40
Instant AP VPN Support

ArubaOS is the companion controller release for the Dell Instant release. This release provides an ability to terminate VPN and GRE tunnels from Instant AP and provide corporate connectivity to the branch Instant AP network. For more details, see the *Dell Networking W-Series Instant Access Point User Guide* .

VPN features are ideal for:

- enterprises with many branches that do not have a dedicated VPN connection to the Head Quarter.
- branch offices that require multiple APs.
- individuals working from home, connecting to the VPN.

This new architecture and form factor seamlessly adds the survivability feature of Instant APs with the VPN connectivity of RAPs — providing corporate connectivity to branches.

This documentation for this feature includes the following topics:

- Overview on page 978
- VPN Configuration on page 982
- Viewing Branch Status on page 983

Overview

This section provides a brief summary of the new features included in ArubaOS to support VPN termination from Instant AP.

Improved DHCP Pool Management

Instant AP (IAP) allows you to configure the DHCP address assignment for the branches connected to the corporate network through VPN. In distributed DHCP mode, ArubaOS 6.3 allows designated blocks of IP addresses for static IP users by excluding them from the DHCP scope. In addition, it allows creation of scope of any required size, thereby enabling more efficient utilization of IP address across branches. For detailed information on Distributed DHCP for IAP-VPN, see *Dell Networking W-Series Instant Access Point User Guide*.

Termination of Instant AP VPN Tunnels

Instant AP (IAP) has the ability to terminate VPN tunnels on controllers. The IAP cluster creates a tunnel from the Virtual Controller to aDell mobility controller in your corporate office. The controller only acts as a VPN end-point and does not configure the IAP. For more information on how to create a VPN tunnel from Virtual Controller to aDell mobility controller, see *Dell Networking W-Series Instant Access Point User Guide*.

Termination of IAP GRE Tunnels

IAPs have the ability to terminate GRE tunnels on controllers. The IAP cluster creates a tunnel from the Virtual Controller to the mobility controller in your corporate office. The controller only acts as a GRE end-point and does not configure the IAP. For more information on how to create a GRE tunnel from Virtual Controller to the mobility controller, see the *Dell Networking W-Series Instant Access Point User Guide*.

L2/L3 Network Mode Support

The Virtual Controller on an Instant AP enables different DHCP pools (various deployment models) in addition to allocating IP subnets to each branch. The following modes of DHCP server are supported:

- L2 Switching Mode: In this mode, Instant supports distributed L2 and centralized L2 switching modes of connection to corporate. When an Instant AP registers with the controller and has a L2 mode DHCP pool configured, the controller automatically adds the GRE or VPN tunnel associated to this IAP into the VLAN multicast table. This allows the clients connecting to this L2 mode VLAN to be part of the same L2 domain on controller.
- L3 Routing Mode: In this mode, Instant supports L3 routing mode of connection to corporate. The VC assigns an IP addresses from the configured subnet and forwards traffic to both corporate and non-corporate destinations. Instant AP takes care of routing on the subnet and also adds a route on the controller after the VPN tunnel is set up during the registration of the subnet. When the Instant AP registers with a L3 mode DHCP pool, the controller automatically adds a route to this DHCP subnet enabling routing of traffic from the corporate to clients on this VLAN in the branch.

Instant AP VPN Scalability Limits

ArubaOS provides enhancements to the scalability limits for the IAP VPN branches terminating on the controller. The following table provides the IAP VPN scalability information for various controller platforms:

Table 217: *Instant AP VPN Scalability Limits*

Platforms	Branches	Routes	L3 Mode Users	NAT Users	Total L2 Users
W-3200	1000	1000	N/A	N/A	64000
W-3400	2000	2000			64000
W-3600	8000	8000			64000
W-6000M3	8000	8000			64000
W-7210	8000	8000			64000
W-7220	16000	16000			128000
W-7240	32000	32000			128000

- **Branches**—The number of IAP VPN branches that can be terminated on a given controller platform.
- **Routes**—The number of L3 routes supported on the controller.
- **L3 mode and NAT mode users**—The number of trusted users supported on the controller. There is no scale impact on the controller. They are limited only by the number of clients supported per Instant AP.
- **L2 mode users**—The number of L2 mode users are limited to128000 for W-7220 and W-7240 and 64000 across all other platforms.

Instant AP VPN OSPF Scaling

ArubaOS allows each IAP VPN to define a separate subnet derived from a corporate intranet pool to allow IAP VPN devices to work independently. For information on sample topology and configuration, see OSPFv2.

To redistribute IAP VPN routes into the OSPF proces, use the following command :

```
(host)(config) # router ospf redistribute rapng-vpn
```

To verify if the redistribution of the IAP VPN is enabled, use following command:

```
(host) #show ip ospf redistribute
Redistribute RAPNG
```

To configure aggregate route for IAP VPN routes, use the following command:

```
(host) (config) # router ospf aggregate-route rapng-vpn
```

To view the aggregated routes for IAP VPN routes, use the following command:

```
(host) #show ip ospf rapng-vpn aggregate-routes
RAPNG VPN aggregate routes
--------------------------
Prefix Mask Contributing routes Cost
------ ---- ------------------- ----
201.201.200.0 255.255.252.0 5 268779624
100.100.2.0 255.255.255.0 1 10
```

To verify the details of configured aggregated route, use the following command:

```
(host) # show ip ospf rapng-vpn aggregated-routes <net> <mask>
(host) #show ip ospf rapng-vpn aggregate-routes 100.100.2.0 255.255.255.0
Contributing routes of RAPNG VPN aggregate route
------------------------------------------------
Prefix Mask Next-Hop Cost
------ ---- -------- ----
100.100.2.64 255.255.255.224 5.5.0.10 10
```

To view all the redistributed routes:

```
(host) #show ip ospf database
OSPF Database Table
-------------------
Area ID    LSA Type      Link ID        Adv Router     Age   Seq#         Checksum
-------    --------      -------        ----------     ---   ----         --------
0.0.0.15   ROUTER        9.9.9.9        9.9.9.9        159   0x80000016   0xee92
0.0.0.15   ROUTER        10.15.148.12   10.15.148.12   166   0x80000016   0x4c0d
0.0.0.15   NETWORK       10.15.148.12   10.15.148.12   167   0x80000001   0x9674
0.0.0.15   NSSA          12.12.2.0      9.9.9.9        29    0x80000003   0x7b54
0.0.0.15   NSSA          12.12.12.0     9.9.9.9        164   0x80000008   0x63a
0.0.0.15   NSSA          12.12.12.32    9.9.9.9        164   0x80000008   0x7b8
0.0.0.15   NSSA          50.40.40.0     9.9.9.9        164   0x80000007   0x8ed4
0.0.0.15   NSSA          51.41.41.128   9.9.9.9        164   0x80000007   0x68f6
0.0.0.15   NSSA          53.43.43.32    9.9.9.9        164   0x80000007   0x2633
0.0.0.15   NSSA          54.44.44.16    9.9.9.9        164   0x80000007   0x353
N/A        AS_EXTERNAL   12.12.2.0      9.9.9.9        29    0x80000003   0x8c06
N/A        AS_EXTERNAL   12.12.12.0     9.9.9.9        169   0x80000001   0x25e4
N/A        AS_EXTERNAL   12.12.12.32    9.9.9.9        169   0x80000001   0x2663
N/A        AS_EXTERNAL   50.40.40.0     9.9.9.9        169   0x80000001   0xab80
N/A        AS_EXTERNAL   51.41.41.128   9.9.9.9        169   0x80000001   0x85a2
N/A        AS_EXTERNAL   53.43.43.32    9.9.9.9        169   0x80000001   0x43de
N/A        AS_EXTERNAL   54.44.44.16    9.9.9.9        169   0x80000001   0x20fe
```

To verify if the redistributed routes are installed or not.

```
(host) #show ip route
Codes: C - connected, O - OSPF, R - RIP, S - static
M - mgmt, U - route usable, * - candidate default, V - RAPNG VPN
Gateway of last resort is Imported from DHCP to network 0.0.0.0 at cost 10
Gateway of last resort is Imported from CELL to network 0.0.0.0 at cost 10
Gateway of last resort is Imported from PPPOE to network 0.0.0.0 at cost 10
Gateway of last resort is 10.15.148.254 to network 0.0.0.0 at cost 1
S*    0.0.0.0/0  [1/0] via 10.15.148.254*
V    12.12.2.0/24 [10/0] ipsec map
V    12.12.12.0/25 [10/0] ipsec map
V    12.12.12.32/27 [10/0] ipsec map
```

```
V    50.40.40.0/24 [10/0] ipsec map
V    51.41.41.128/25 [10/0] ipsec map
V    53.43.43.32/27 [10/0] ipsec map
V    54.44.44.16/28 [10/0] ipsec map
C    9.9.9.0/24 is directly connected, VLAN9
C    10.15.148.0/24 is directly connected, VLAN1
C    43.43.43.0/24 is directly connected, VLAN132
C    42.42.42.0/24 is directly connected, VLAN123
C    44.44.44.0/24 is directly connected, VLAN125
C    182.82.82.12/32 is an ipsec map 10.15.149.69-182.82.82.12
C    182.82.82.14/32 is an ipsec map 10.17.87.126-182.82.82.14
```

Branch-ID Allocation

For branches deployed in distributed L3 and distributed L2 mode, the master AP in the branch and the controller should agree upon a subnet/IP addresses to be used for DHCP services in the branch. The process or protocol used by the master AP and the controller to determine the subnet/IP addresses used in a branch is called BID allocation. The BID allocation process is not essential for branches deployed in local or centralized L2 mode. The following are some of the key functions of the BID allocation process:

- Determines the IP addresses used in a branch for distributed L2 mode
- Determines the subnet used in a branch for distributed L3 mode
- Avoids IP address or subnet overlap (that is, avoids IP conflict)
- Ensures that a branch is allocated the same subnet or range of IP addresses irrespective of which AP in the branch becomes the master in the IAP cluster

Centralized BID Allocation

In Master-Local controller set-up, the Master controller runs the BID allocation algorithm and allocates BID to the branches terminating on it and to the Local controllers. The Master controller saves the BIDs in it's memory IAP database to avoid the collision of BID (per subnet) whereas the Local controller saves the BIDs only in it's in memory data structures. The IAP manager in Local controller forwards only the new register request (branch coming for the first time with BIDs as -1) message to the Master controller. For existing branch's register request, the Local controller tries to honor the requested BIDs first. The master and local communication is within the existing IPSec tunnel. The Master controller gets the register request and allocates BIDs using the BID allocation algorithm. Finally, the Master controller sends back the allocated BIDs to the Local controller and the Local controller updates its data structure and sends the response to the IAP.

General guidelines for upgrading from existing IAP-VPN release to ArubaOS 6.4:

1. Ensure that all the branches are upgraded to Instant 4.0.
2. Upgrade the data-center to ArubaOS 6.4.



If you have a Master-Local setup; upgrade the Master controller first and then the Local controller.

3. Ensure that always the IAP-VPN branches are configured using authorized tools like AirWave, else you must trust all branches or the required branch using the following command,

   ```
   iap trusted-branch-db allow-all
   ```

   or

   ```
   iap trusted-branch-db add mac-address<mac-address>
   ```



Instant version earlier than 4.0 also need the previous command to be executed for the controller to come up with ArubaOS 6.4.

VPN Configuration

The following VPN configuration steps on the controller, enable IAPs to terminate their VPN connection on the controller:

Whitelist DB Configuration

Controller Whitelist DB

You can use the following CLI command to configure the whitelist DB if the controller is acting as the whitelist entry:

```
(host) #whitelist-db rap add mac-address 00:11:22:33:44:55 ap-group test
```

The **ap-group** parameter is not used for any configuration, but needs to be configured. The parameter can be any valid string. If an external whitelist is being used, the MAC address of the AP needs to be saved in the Radius server as a lower case entry without any delimiter.

External Whitelist DB

The external whitelist functionality enables you to configure the RADIUS server to use an external whitelist for authentication of MAC addresses of RAPs.

If you are using Windows 2003 server, perform the following steps to configure external whitelist on it. There are equivalent steps available for Windows Server 2008 and other RADIUS servers.

1. Add the MAC addresses for all the RAPs in the Active Directory of the Radius server:
   a. Open the **Active Directory and Computers** window, add a new user and specify the MAC address (without the colon delimiter) of the RAP for the user name and password.
   b. Right-click the user that you have just created and click **Properties**.
   c. In the **Dial-in** tab, select **Allow access** in the **Remote Access Permission** section and click **OK**.
   d. Repeat Step a through Step b for all RAPs.
2. Define the remote access policy in the Internet Authentication Service:
   a. In the **Internet Authentication Service** window, select **Remote Access Policies**.
   b. Launch the wizard to configure a new remote access policy.
   c. Define filters and select select **grant remote access permission** in the **Permissions** window.
   d. Right-click the policy that you have just created and select **Properties**.
   e. In the **Settings** tab, select the policy condition, and **Edit Profile...**.
   f. In the **Advanced** tab, select **Vendor Specific**, and click **Add** to add new vendor specific attributes.
   g. Add new vendor specific attributes and click **OK**.
   h. In the **IP** tab, provide the IP address of the RAP and click **OK**.

VPN Local Pool Configuration

The VPN local pool is used to assign an IP Address to the IAP after successful XAUTH VPN.

```
(host) # ip local pool "rapngpool" <startip> <endip>
```

Role Assignment for the Authenticated IAPs

Define a role that includes a source NAT rule to allow connections to the RADIUS server and for the Dynamic Radius Proxy in the IAP to work. This role is assigned to IAPs after successful authentication.

```
(host) (config) #ip access-list session iaprole
(host) (config-sess-iaprole) #any host <radius-server-ip> any src-nat
(host) (config-sess-iaprole) #any any any permit
(host) (config-sess-iaprole) #!
```

```
(host) (config) #user-role iaprole
(host) (config-role) #session-acl iaprole
```

VPN Profile Configuration

The VPN profile configuration defines the server used to authenticate the IAP (internal or an external server) and the role assigned to the IAP after successful authentication.

```
(host) (config) #aaa authentication vpn default-iap
(host) (VPN Authentication Profile "default-iap") #server-group default
(host) (VPN Authentication Profile "default-iap") #default-role iaprole
```



The **default role** parameter of the **aaa authentication vpn** command requires **Policy Enforcement Firewall for VPN users (PEFV)** license.



By default, the controller uses the default IAP role. If the administrator changes the IAP role name when the IAP's status is UP, then the controller or the IAP must be rebooted.

For more information on VPN profile configuration, see the *VPN Configuration* chapter of the *Dell Networking W-Series Instant Access Point User Guide*.

Viewing Branch Status

To view the details of the branch information connected to the controller, execute the **show iap table** command.

Example

This example shows the details of the branches connected to the controller:

```
(host) #show iap table long

IAP Branch Table
----------------
Name             VC MAC Address     Status   Inner IP       Assigned Subnet  Assigned Vlan
----             --------------     ------   --------       ---------------  -------------
Tokyo-CB:D3:16   6c:f3:7f:cc:42:f8  DOWN     0.0.0.0
Paris-CB:D3:16   6c:f3:7f:cc:3d:04  UP       10.15.207.140  10.15.206.99/29  2
LA               6c:f3:7f:cc:42:25  UP       10.15.207.111  10.15.206.24/29  2
Munich           d8:c7:c8:cb:d3:16  DOWN     0.0.0.0
London-c0:e1     6c:f3:7f:c0:e1:b1  UP       10.15.207.120  10.15.206.64/29  2
Instant-CB:D3    6c:f3:7f:cc:42:1e  DOWN     0.0.0.0
Delhi            6c:f3:7f:cc:42:ca  DOWN     0.0.0.0
Singapore        6c:f3:7f:cc:42:cb  UP       10.15.207.122  10.15.206.120/29 2


Key        Bid(Subnet Name)
---        ----------------
b3c65c...
b3c65c...
b3c65c...  2(10.15.205.0-10.15.205.250,5),1(10.15.206.1-10.15.206.252,5)
a2a65c...  0
b3c65c...  7(10.15.205.0-10.15.205.250,5),8(10.15.206.1-10.15.206.252,5)
b3c65c...
b3c65c...  1(10.15.205.0-10.15.205.250,5),2(10.15.206.1-10.15.206.252,5)
b3c65c...  14(10.15.205.0-10.15.205.250,5),15(10.15.206.1-10.15.206.252,5)
```

The output of this command includes the following parameters:

Table 218: *IAP Table Parameters*

Parameter	Description
`Name`	Displays the name of the branch.
`VC MAC Address`	Displays the MAC address of the Virtual Controller of the branch.
`Status`	Displays the current status of the branch (UP/DOWN).
`Inner IP`	Displays the internal VPN IP of the branch.
`Assigned Subnet`	Displays the subnet mask assigned to the branch.
`Assigned Vlan`	Displays the VLAN ID assigned to the branch.
`Key`	Displays the key for the branch, which is unique to each branch.
`Bid(Subnet Name)`	Displays the Branch ID (BID) of the subnet. • In the example above, the controller displays bid-per-subnet-per-branch i.e., for "LA" branch, BID "2" for the ip-range "10.15.205.0-10.15.205.250" with client count per branch "5"). If a branch has multiple subnets, it can have multiple BIDs. • Branches that are in **UP** state and do not have a **Bid(Subnet Name)** means that the IAP is connected to a controller which did not assign any bid for any subnet. In the above example, "Paris-CB:D3:16" branch is **UP** and does not have a **Bid (Subnet Name)** information. This means that either the IAP is connected to a backup controller or connected to a primary controller without any distributed L2 or L3 subnets. For more information on bid-per-subnet-per-branch and distributed L2 and L3 subnets, see the *DHCP Configuration* chapter of the Dell *Instant Access Point 6.2.1.0-3.3 User Guide*.



Executing the **show iap table** command does not display the **Key** and **Bid(Subnet Name)** parameters.

The W-600 Series Controller is designed for compact, cost-effective "all-in-one" networking solutions. The W-600 Series includes a firewall, wireless LAN controller, Ethernet switch with PoE+, IP router, site-to-site VPN edge device, file server, and print server.

The W-600 Series is an enterprise-class, wireless LAN controller that connects, controls, and integrates wireless APs and Air Monitors (AMs) into a wired LAN system. Table 219 lists some of the hardware features by the numbers.

Table 219: *W-600 Series Controller by the Numbers*

Controller	USB Ports	Maximum External APs	Remote APs
W-620	1	8	8
W-650	4	16	16

Topics in this chapter include:

- Connecting with a USB Cellular Modems on page 985
- Configuring a Supported USB Modem on page 987
- Configuring a New USB Modem on page 988

Connecting with a USB Cellular Modems

USB Cellular Modems are supported via a USB port. ArubaOS supports several EVDO (Evolution Data Optimized, up to 3.1 Mbps, CDMA) and 3G HSPA (High-Speed Packet Access, 3G data service), and 4G LTE (Fourth Generation, Long Term Evolution) modems. The 3G HSPA is provided by AT&T in the United States and numerous other 3G providers worldwide.

How it Works

Plug the USB Cellular Modem into the USB port of the W-600 Series controller. The USB Cellular Modem is automatically detected and negotiates a PPP IP address. If the modem fails to obtain a PPP IP address within 45 seconds, the controller ignores the modem's presence, and boots as if the modem is not present.

Switching Modes

Many of the newer modems contain multiple USB devices; creating a very elegant plug-n-play solution. When your USB Cellular Modem is first powered on, a storage device is registered. This storage device contains the software driver/executable necessary to install and operate the modem.

Once the software installation is complete, the modem must *mode-switch* from a storage device to a registered modem device. Mode-switching varies by manufacturer. For example, The Novatel modem mode-switches via a SCSI eject command; the Huawei modem mode-switches via a SCSI rezero command, while the Sierra modem mode-switches via a specific USB command. Once the mode-switching is complete, the modem automatically registers itself.

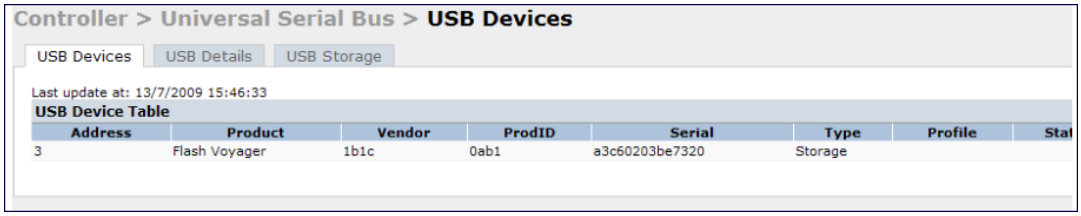
The controller can dial (via the modem) your Service Provider to initiate a PPP session. During the boot sequence, the controller issues your device's mode-switching command, every few seconds, until the PPP link connects.

Finding USB Modem Commands

To support the USB cellular modems on the W-600 Series, the **cellular profile <profile>** and **uplink** commands are available at the command line interface. For detailed information on these commands, refer to the Command Line Reference Guide.

You can view connected USB cellular devices via the **Controller > Universal Serial Bus > USB Devices** in the Web UI (see Figure 224). Navigating to this page is the equivalent of executing the `show usb` command at the command prompt.

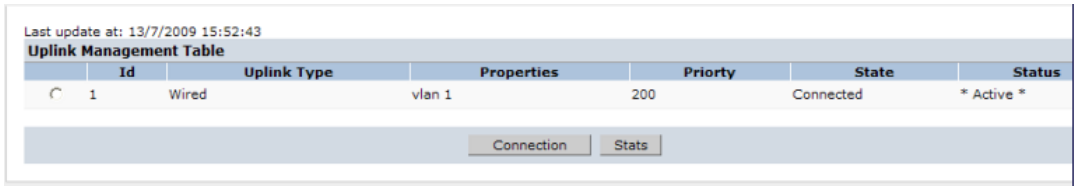
Figure 224 *Connected Cellular Devices*



Uplink Manager

Access the Uplink Manager feature from the WebUI Configuration tab. Navigate to this feature via **Uplink > Uplink Manager** ().

Figure 225 *WebUI Uplink Manager*



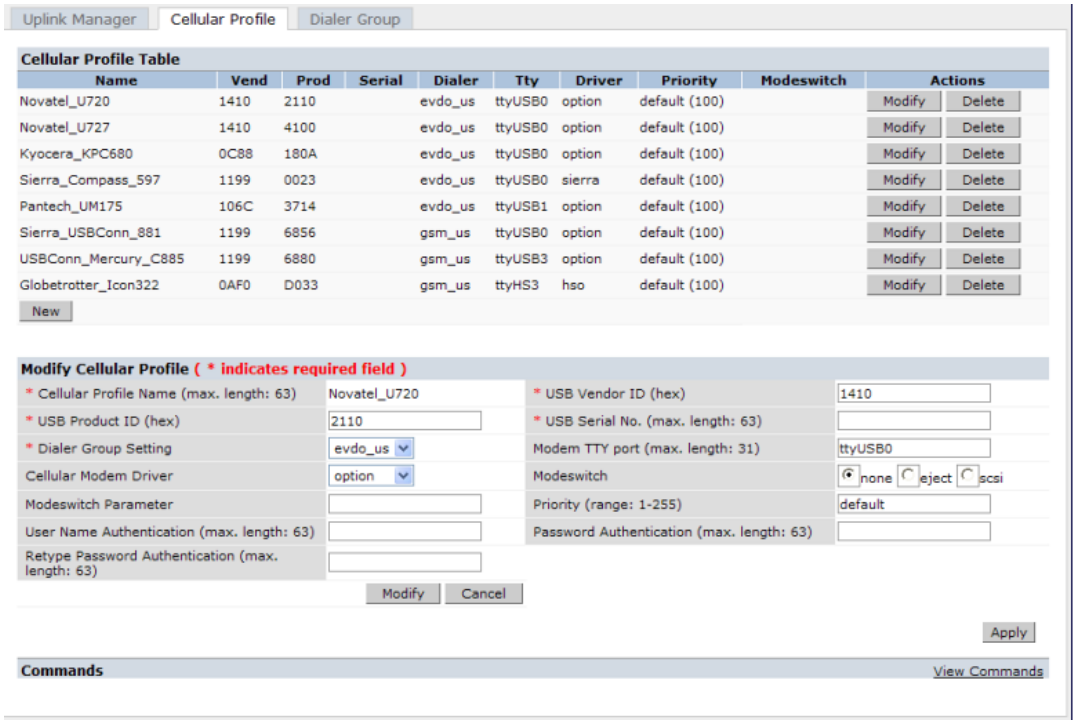
You can enable/disable the uplink to overwrite cellular and wired uplink priority. The corresponding commands are:

```
(host) (config)# uplink [enable | disable]
(host) (config)# uplink [cellular | wired] priority [x]
```

Cellular Profile

The Cellular Profile tab allows you to add/modify/delete one or more cellular profiles. The WebUI screen for Cellular Profile is divided into the Cellular Profile Table (the top portion) and the Modify Cellular Profile (the bottom portion). When a cellular profile is selected for modification (see Figure 226) the bottom modify portion is revealed. All changes are entered into the buffer until the Apply button is executed.

Figure 226 *Cellular Profile from the WebUI*

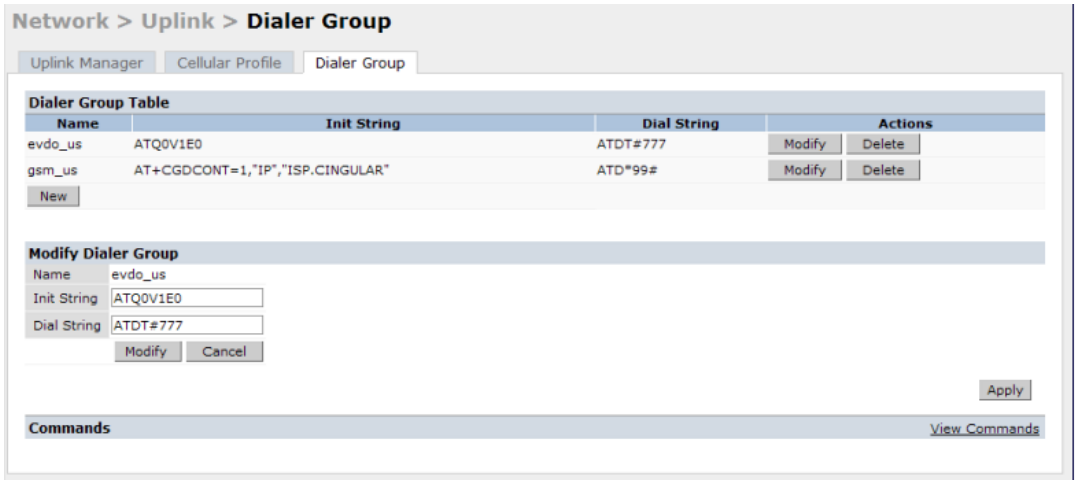


Dialer Group

Use the Dialer Group command to configure EVDO devices that require specific input for the initial string (init-string) and dial string. When adding or modifying an existing dialer group (see Figure 227), the WebUI executes the following commands:

```
(host) (config-cellular profile_name)# dialer group <name> init-string <string>
(host) (config-cellular profile_name)# dialer group <name> dial-string <string>
```

Figure 227 *Configuring Dialer Group*



Configuring a Supported USB Modem

If your USB Modem is a validated modem, then no configuration is needed. Just follow the “plug and play” steps below.

1. Insert the USB Modem into an open USB port.

2. From the controller command-line interface, issue the show usb command to that the modem is detected.

```
(host) (config-cellular new_modem)# show usb
USB Device Table
----------------
Address  Product                Vendor  ProdID  Serial           Type      Profile    State
-------  -------                ------  ------  ------           ----      -------    -----
18       Novatel Wireless CDMA  1410    4100    091087843891000  Cellular  new_modem
Device ready
```

3. If your modem is not recognized (such as “type is unknown”, “no matching profile”, or “device not ready”), use the **`show usb verbose`** command to verify your modem is listed.

```
(host) #show usb verbose
...
T: Bus=01 Lev=02 Prnt=02 Port=00 Cnt=01 Dev#= 3 Spd=12 MxCh= 0
D: Ver= 1.10 Cls=00(>ifc ) Sub=00 Prot=00 MxPS=64 #Cfgs= 1
P: Vendor=1410 ProdID=4100 Rev= 0.00
S: Manufacturer=Novatel Wireless Inc.
S: Product=Novatel Wireless CDMA
S: SerialNumber=091087843891000
C:* #Ifs= 5 Cfg#= 1 Atr=a0 MxPwr=500mA
...
```

4. Issue the show uplink command to verify the modem as an active status, and is registered with the Uplink Manager.

   Cellular uplinks have a lower priority than wired links by default. You can change the default by changing the profile-specific priority or by changing the default cell priority.

```
(host) (config) #cellular priority <prior>
```

5. Check the modem dialing status. The connection may take up to a 45 seconds to establish. To see the connection progress, execute the **show uplink connection uplink id** command.
6. Verify the connection is established and IP addressed is programmed.
   - Once the cellular link state is *Connected*, you can find the PPP dynamic entries by executing the command **`show uplink connection id`**
   - The IP address can be found using the command **`show ip interface brief`**
   - The Gateway can be found using the command **`show ip route`**
   - The DNS entries can be found using the command **`show ip domain-name`**

Configuring a New USB Modem

Cellular modems must be activated before they can “talk” on the cellular network. Typically, the activation is done by the carrier. Some carriers use a proprietary PC client. In all cases, make sure that your modem works on your PC before using it on the W-600 Series.



Verify your modem is activated and works with your Microsoft Windows or Apple Mac computers.

Each time a USB device is inserted, Linux assigns it a new USB address. This is true even if the same device is re-inserted. Modem ports are organized under their individual addresses. For example, **ttyUSB0** at address 3 is separate than **ttyUSB0** at address 7. The address is displayed when you execute the commands, **`show usb`** and **`show usb verbose`** **(**the Dev# field).

Configuring the Profile and Modem Driver

1. Insert the USB Modem into an open USB port.
2. Verify that the modem is detected (the **`show usb`** command. (see Figure 233.)
3. If your modem is not recognized (such as “type is unknown”, “no matching profile”, or “device not ready”), use the **`show usb verbose`** (see Figure) command to verify your modem is listed.
4. From the command-line interface, issue the cellular profile <profile> command to reate a cellular profile and import the identifiers. The Dialer, TTY, and Driver fields are the new profile defaults
5. Configure the modem driver.

   The default “option” driver is a catch-all for cellular modems. Nearly all cards use this driver and support for new modems are added here. Once option driver is configured to work with this device, it recognizes the modem and expose its ports.

   If you get entries similar to the example below:

Figure 228 *Driver=(none)*

```
(host) #show usb verbose
...
I:  If#= 0 Alt= 0 #EPs= 3 Cls=ff(vend.) Sub=ff Prot=ff Driver=(none)
I:  If#= 1 Alt= 0 #EPs= 2 Cls=ff(vend.) Sub=ff Prot=ff Driver=(none)
```

This means the driver does not work with these ports. Try the other drivers and see if they pick up the device. Airprime is the reliable *catch-all* driver, Sierra is for certain Sierra cards, and cdc-acm is a legacy abstract control modem driver. Your goal is to assign a driver for the unclaimed (none) interfaces (If#).

If no option driver appears or only storage interfaces appear, then the modem must be switched to data mode (see Switching Modes on page 985).

Configuring the TTY Port

1. View the exposed TTY ports by executing the **show usb ports 13** command.

Figure 229 *show usb ports 13 command*

```
(host) (config-cellular new_card)# show usb ports 13
ttyUSB0
ttyUSB1
ttyUSB2
ttyUSB3
```

In the example above, the command reveals four exposed TTY ports. One is the modem port, while the other ports are for GPS, real-time statistics, or diagnostics. If the command does not reveal any ports or if only storage devices (such as ‘sr0’) appear, then the device must be switched to data mode before proceeding. See Switching Modes on page 985 for instruction.

2. Send a test AT command to determine the correct modem port.

Figure 230 *show usb test command*

```
(host) (support)#show usb test 16 ttyUSB0
AT
OK
TTY port responded to modem AT commands
```

In the example above, the TTY port responds with an ‘OK’. This indicates that ttyUSB0 is a valid modem port.

There may be more than one modem port; you can continue to send AT commands to determine which ports are modem ports. If the port is not a valid modem port, a time out error is generated as shown in the example below

Figure 231 *Time out error example.*

```
(host) (support)#show usb test 16 ttyUSB1
Error: Timed out while waiting for modem to respond to AT commands
(host) (support)#
```

In the example below, the TTY port does not exist, or is busy with a previous PPP session.

Figure 232 *Port I/O error*

```
(host) (support)#show usb test 16 ttyUSB4
Error: Port I/O error. TTY port usb/16/ttyUSB4 inaccessible
(host) (support)#
```

Once you find one (or more) modem TTY port, configure it in the cellular profile and test the port.

Testing the TTY Port

After your TTY port is correctly configured, the port is in the 'Device Ready' state.

Figure 233 *Device Ready State*

```
(host) (config-cellular new_modem)# show usb
USB Device Table
----------------
Address  Product                Vendor  ProdID  Serial           Type      Profile    State
-------  -------                ------  ------  ------           ----      -------    -----
18       Novatel Wireless CDMA  1410    4100    091087843891000  Cellular  new_modem  Deviceready
(host) (config-cellular new_modem)#
```

The 'Device Ready' state indicates the port has passed the diagnostic test and is ready.

You can also run extended diagnostics to displays more information about the modem.



Not all modems support the extended AT command set. If the modem hangs after sending an extended AT command; removing the device and then re-inserting it usually fixes the problem

The AT+CSQ command queries is the modem's current signal strength. The first number represents the signal ranging from 1 (poor) to 33 (excellent). In the example below, the strength is in the excellent range (31).

Figure 234 *usb test extended.*

```
(host) #show usb test 18 ttyUSB0 extended
OK
ATI0
Manufacturer: NOVATEL WIRELESS INCORPORATED
Model: U727 SPRINT
Revision: m6800B-RAPTOR65_S-114 [Dec 07 2007 18:00:00]
ESN: 0x5B860A05
+GCAP: +CIS707-A, CIS-856-A, +MS, +ES, +DS
OK
AT+CSQ
31, 99
```

Selecting the Dialer Profile

The phone number, user name, and password (if any) are set in the dialer setting. In the United States, AT&T and T-Mobile use the 'gsm_us' profile, while Sprint and Verizon use the 'evdo_us' profile. User names and passwords are not typically used by U.S. carriers, but they may be required by International carriers.

Choose the dialer group that matches your carrier. If one doesn't exist, create a new dialer group with information from your carrier (Figure 235)

Figure 235 *show dialer group example*

```
(host)# show dialer group
Dialer Group Table
------------------
Name         Init String                      Dial String
----         -----------                      -----------
evdo_us      ATQ0V1E0                         ATDT#777
gsm_us       AT+CGDCONT=1,"IP","ISP.CINGULAR" ATD*99#
```

The ATD, in the Dial String column in Figure 235, specifies the number to dial, and is typically the same among respective CDMA/GSM carriers. The information under the Init String column typically just resets the modem to the factory default state, but may contain carrier specific options. You can often find these settings in online forums or from your ISP.

Linux Support

The Internet is a great place to research Linux support for your modem. Chances are someone already got it working on their system and their configuration can be leveraged. The following sites provide useful information:

- evdoforums.com/
- ubuntuforums.org
- linux.com/forums
- kenkinder.com/

The External Services Interface (ESI) provides an open interface that is used to integrate security solutions that solve interior network problems such as viruses, worms, spyware, and corporate compliance. ESI allows selective redirection of traffic to external service appliances such as anti-virus gateways, content filters, and intrusion detection systems. When "interesting" traffic is detected by these external devices, it can be dropped, logged, modified, or transformed according to the rules of the device. ESI also permits configuration of different server groups—with each group potentially performing a different action on the traffic.

You can configure ESI to do one or more of the following for each group:

- Redirect specified types of traffic to the server
- Perform health checks on each of the servers in the group
- Perform per-session load balancing between the servers in each group
- Provide an interface for the server to return information about the client that can place the client in special roles such as "quarantine"

ESI also provides the ESI syslog parser, which is a mechanism for interpreting syslog messages from third-party appliances such as anti-virus gateways, content filters, and intrusion detection systems. The ESI syslog parser is a generic syslog parser that accepts syslog messages from external devices, processes them according to user-defined rules, and then takes configurable actions on system users.

Topics in this chapter include:

- Sample ESI Topology on page 992
- Understanding the ESI Syslog Parser on page 994
- Configuring ESI on page 996
- Sample Route-mode ESI Topology on page 1004
- Sample NAT-mode ESI Topology on page 1010
- Understanding Basic Regular Expression (BRE) Syntax on page 1016

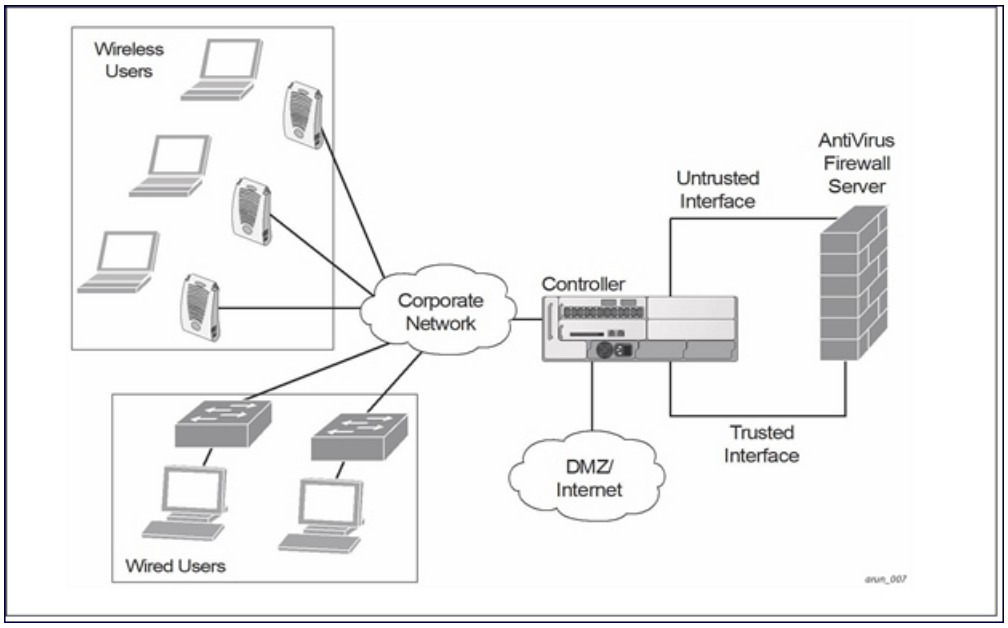


The ESI feature requires the Policy Enforcement Firewall Next Generation (PEFNG) license installed on the controller.

Sample ESI Topology

In the example shown in this section, ESI is used to provide an interface to the AntiVirusFirewall (AVF) server device for providing virus inspection services. An AVF server device is one of many different types of services supported in the ESI.

Figure 236 *ESI-Fortinet Topology*



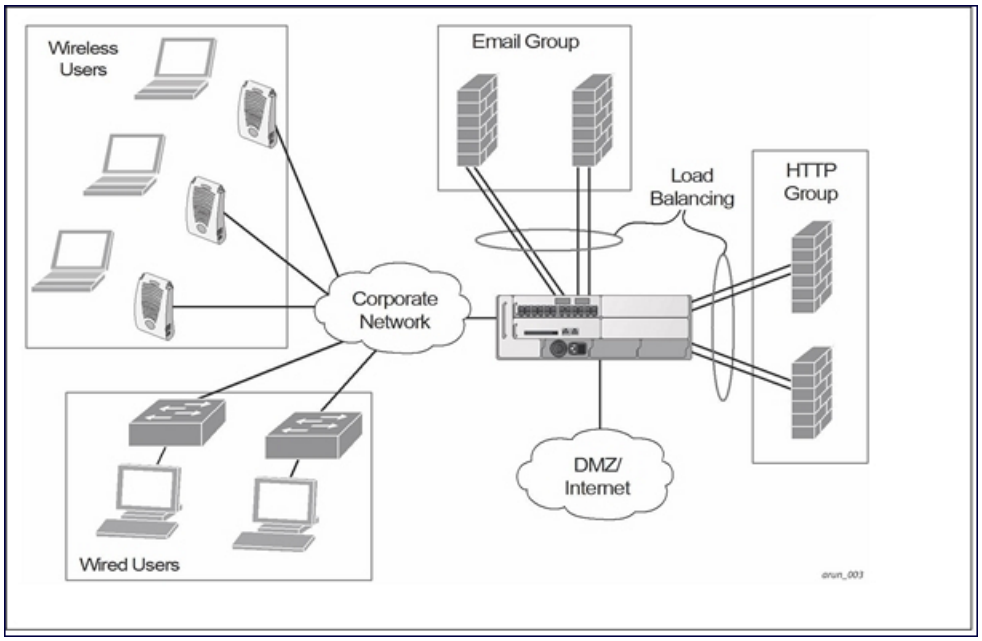
In the ESI–Fortnet topology, the clients connect to access points (both wireless and wired). The wired access points tunnel all traffic back to the controller over the existing network.

The controller receives the traffic and redirects relevant traffic (including but not limited to all HTTP/HTTPS and email protocols such as SMTP and POP3) to the AVF server device to provide services such as anti-virus scanning, email scanning, web content inspection, etc. This traffic is redirected on the “untrusted” interface between the controller and the AVF server device. The controller also redirects the traffic intended for the clients coming from either the Internet or the internal network. This traffic is redirected on the “trusted” interface between the controller and the AVF server device. The controller forwards all other traffic (for which the AVF server does not perform any of the required operations such as AV scanning). An example of such traffic would be database traffic running from a client to an internal server.

The controller can also be configured to redirect traffic only from clients in a particular role such as “guest” or “non-remediated client” to the AVF server device. This might be done to reduce the load on the AVF server device if there is a different mechanism such as the Dell-Sygate integrated solution to enforce client policies on the clients that are under the control of the IT department. These policies can be used to ensure that an anti-virus agent runs on the clients and the client can get access to the network only if this agent reports a “healthy” status for the client. Refer to the paper (available from Sygate) on Sygate integrated solutions for more details on this solution.

The controller is also capable of load balancing between multiple external server appliances. This provides more scalability as well as redundancy by using multiple external server appliances. Also, the controller can be configured to have multiple groups of external server devices and different kinds of traffic can be redirected to different groups of devices with load balancing occurring within each group (see Figure 237 for an example).

Figure 237 *Load Balancing Groups*



Understanding the ESI Syslog Parser

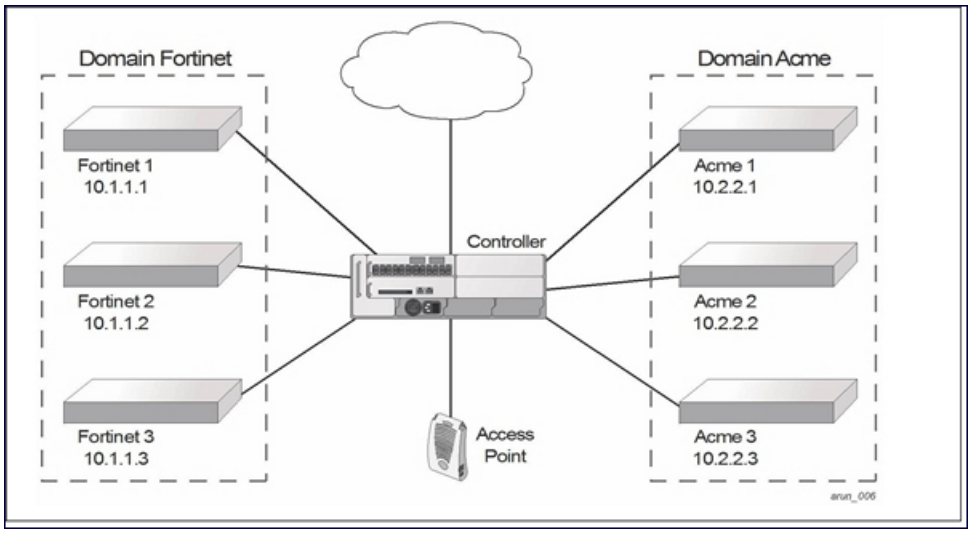
The ESI syslog parser adds a UNIX-style regular expression engine for parsing relevant fields in messages from third-party appliances such as anti-virus gateways, content filters, and intrusion detection systems.

The user creates a list of rules that identify the type of message, the username to which this message pertains, and the action to be taken when there is a match on the condition.

ESI Parser Domains

The ESI servers are configured into ESI parser domains (see Figure 238) to which the rules will be applied. This condition ensures that only messages coming from configured ESI parser domains are accepted, and reduces the number of rules that must be examined before a match is detected (Syslog Parser Rules on page 996). messages.When a syslog message is received, it is checked against the list of defined ESI servers. If a server match is found, the message is then tested against the list of predefined rules.

Figure 238 *ESI Parser Domains*



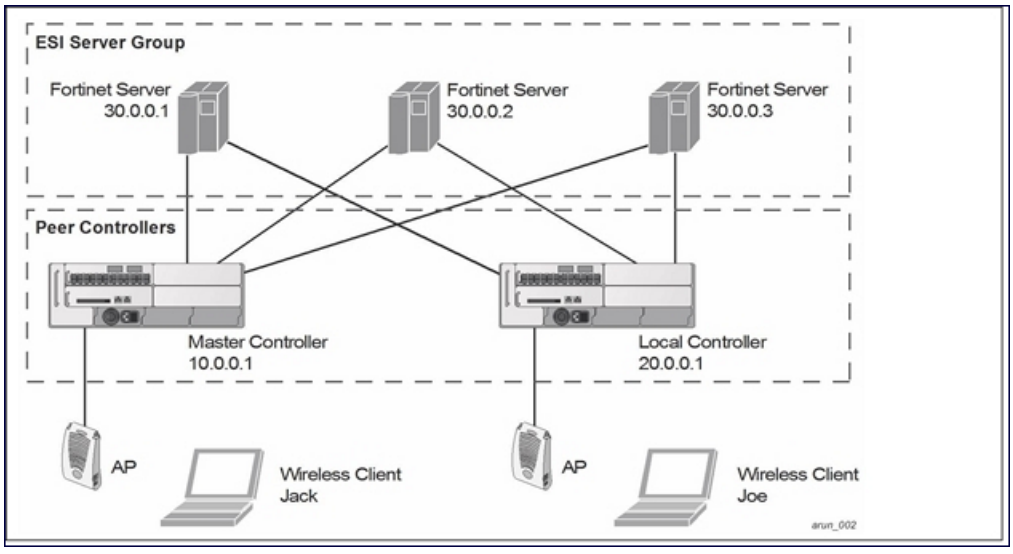
The ESI syslog parser begins with a list of configured IP interfaces which listen for ESI messages. When a syslog message is received, it is checked against the list of defined ESI servers. If a server match is found, the message is then tested against the list of predefined rules.

Within the rule-checking process, the incoming message is checked against the list of rules to search first for a condition match (see Syslog Parser Rules on page 996). If a condition match is made, and the user name can be extracted from the syslog message, the resulting user action is processed by first attempting to look up the user in the local user table. If the user is found, the appropriate action is taken on the user. The default behavior is to look for users only on the local controller. If the user is not found, the event is meaningless and is ignored. This is the typical situation when a single controller is connected to a dedicated ESI server.

Peer Controllers

As an alternative, consider a topology where multiple controllers share one or more ESI servers.

Figure 239 *ESI Peer Controllers*



In this scenario, several controllers (master and local) are defined in the same syslog parser domain to act as *peers*. From the standpoint of the ESI servers, because there is no accurate way of determining from which controller a given user came. Thus, the event is flooded out to all controllers defined as peers within this ESI

parser domain. The corresponding controller holding the user entry acts on the event, while other controllers ignore the event.

Syslog Parser Rules

The user creates an ESI rule by using characters and special operators to specify a pattern (regular expression) that uniquely identifies a certain amount of text within a syslog message. (Regular expression syntax is described in Understanding Basic Regular Expression (BRE) Syntax on page 1016.) This "condition" defines the type of message and the ESI domain to which this message pertains. The rule contains three major fields:

- Condition: The pattern that uniquely identifies the syslog message type.
- User: The username identifier. It can be in the form of a name, MAC address, or IP address.
- Action: The action to take when a rule match occurs.

Once a condition match has been made, no further rule-matching will be made. For the rule that matched, only one action can be defined.

After a condition match has been made, the message is parsed for the user information. This is done by specifying the target region with the regular expression (REGEX) regex() block syntax. This syntax generates two blocks: The first block is the matched expression; the second block contains the value inside the parentheses. For username matching, the focus is on the second block, as it contains the username.

Condition Pattern Matching

The following description uses the Fortigate virus syslog message format as an example to describe condition pattern matching. The Fortigate virus syslog message takes the form:

```
Sep 26 18:30:02 log_id=0100030101 type=virus subtype=infected src=1.2.3.4
```

This message example contains the Fortigate virus log ID number 0100030101 ("log_id=0100030101"), which can be used as the condition—the pattern that uniquely identifies this syslog message.

The parser expression that matches this condition is "log_id=0100030101". This is a narrow match on the specific log ID number shown in the message, or "log_id=[0–9]{10}[ ]" ,which is a regular expression that matches any Fortigate log entry with a ten-digit log ID followed by a space.

User Pattern Matching

To extract the user identifier in the example Fortigate virus message shown above ("src=1.2.3.4"), use the following expression, "src=(.*)[ ]" to parse the user information contained between the parentheses. The () block specifies where the username will be extracted. Only the first block will be processed.

More examples:

Given a message wherein the username is a MAC address:

```
Sep 26 18:30:02 log_id=0100030101 type=virus subtype=infected mac 00:aa:bb:cc:dd:00
```

The expression "mac[ ](.{17})" will match "mac 00:aa:bb:cc:dd:00" in the example message.

Given a message wherein the username is a user name:

```
Sep 26 18:30:02 log_id=0100030101 type=virus subtype=infected user<johndoe>
```

The expression "user<(.*)>" will match "user<johndoe>" in the example message.

Configuring ESI

You can use the following interfaces to configure and manage ESI and ESI syslog parser behavior:

- The Web user interface (WebUI), which is accessible through a standard Web browser from a remote management console or workstation.
- The command line interface (CLI), which is accessible from a local console device connected to the serial port on the controller or through a Telnet or Secure Shell (SSH) connection from a remote management console or workstation.



By default, you can access the CLI only from the serial port or from an SSH session. To use the CLI in a Telnet session, you must explicitly enable Telnet on the controller. The general configuration descriptions in the following sections include both the WebUI pages and the CLI configuration commands. The configuration overview section is followed by several examples that show specific configuration procedures.

In general, there are three ESI configuration "phases" on the controller as a part of the solution:

- The first phase configures the ESI *ping health-check method*, *servers*, and *server groups*.The term *server* here refers to external server devices, for example, an AVF.
- The second phase configures the redirection policies instructing the controller how to redirect the different types of traffic to different server groups.
- The final phase configures the ESI syslog parser domains and the rules that interpret and act on syslog message contents.



The procedures shown in the following sections are general descriptions. Your application might be broader or narrower than this example, but the same general operations apply.

Configuring Health-Check Method, Groups, and Servers

To configure the ESI health-check method, servers, and server groups, navigate to the **Configuration > Advanced Services > External Services** view on the WebUI.

In the WebUI

To configure a health check profile:

1. Navigate to the **Configuration > Advanced Services > External Services** page on the WebUI.
2. Click **Add** in the **Health Check Configuration** section.
   (To change an existing profile, click **Edit**.)
3. Provide the following details:
   a. Enter a **Profile Name**.
   b. **Frequency (secs)**—Indicates how often the controller checks to see if the server is up and running. Default: 5 seconds.
   c. **Timeout (secs)**—Indicates the number of seconds the controller waits for a response to its health check query before marking the health check as failed. Default: 2 seconds.
   d. **Retry count**—Is the number of failed health checks after which the controller marks the server as being down. Default: 2.
4. Click **Done** when you are finished.
5. Click **Apply** to apply the configuration changes.

In the CLI

Use these CLI commands to configure a health-check method:

```
esi ping profile_name
   frequency seconds
   retry-count count
   timeout seconds
```

For example:

```
esi ping default
   frequency 5
   retry-count 2
```

Defining the ESI Server

The following sections describe how to configure an ESI server using the WebUI and CLI.

In the WebUI

To configure an ESI server:

1. Navigate to the **Configuration > Advanced Services > External Services** page on the WebUI.
2. Click **Add** in the **External Servers** section.
3. Provide the following details:
   a. **Server Name**.
   b. **Server Group**. Use the drop-down list to assign this server to a group from the existing configured groups.
   c. **Server Mode**. Use the drop-down list to choose the mode (bridge, nat, or route) your topology requires. Refer to the description above to understand the differences between these modes.

      For **routed** mode, enter the **Trusted IP Address** (the IP address of the trusted interface on the external server device) and the **Untrusted IP Address** (the IP address of the untrusted interface on the external server device). (You can also choose to enable a health check on either or both of these interfaces.)

      For **bridged** mode, enter the **Trusted Port** number (the port connected to the trusted side of the ESI server) and the **Untrusted Port** number (the port connected to the untrusted side of the ESI server).

      For **NAT** mode, enter the **Trusted IP Address** (the trusted interface on the external server) and the **NAT Destination Port** number (the port a packet is redirected to rather than the original destination port in the packet). You can also choose to enable a health check on the trusted IP address interface.
4. Click **Done** when you are finished.
5. Click **Apply** to apply the configuration changes..

In the CLI

Use these CLI commands to configure an ESI server and identify its associated attributes:

```
esi server server_identity
   dport destination_tcp/udp_port
   mode {bridge | nat | route}
   trusted-ip-addr ip-addr [health-check]
   trusted-port slot/port
   untrusted-ip-addr ip-addr [health-check]
   untrusted-port slot/port
```

For example:

```
esi server forti_1
   mode route
   trusted-ip-addr 10.168.172.3
   untrusted-ip-addr 10.168.171.3
```

Defining the ESI Server Group

The following sections describe how to configure an ESI server group using the WebUI and CLI.

In the WebUI

To configure an ESI server group on the controller:

1. Navigate to the **Configuration > Advanced Services > External Services** page.
2. Click **Add** in the **Server Groups** section.

   (To change an existing group, click **Edit**.)
3. Provide the following details:
   a. Enter a **Group Name**.
   b. In the drop-down list, select a health check profile.
4. Click **Done** when you are finished.
5. Click **Apply** to apply the configuration changes.

In the CLI

Use these CLI commands to configure an ESI server group, identify its associated ping health-check method, and associate a server with this group:

```
esi group name
   ping profile_name
   server server_identity
```

For example:

```
esi group fortinet
   ping default
   server forti_1
```

Redirection Policies and User Role

The following sections describe how to configure the redirection policies and user role using the WebUI and CLI.

In the WebUI

To configure user roles to redirect the required traffic to the server(s), navigate to the **Configuration > Access Control > User Roles** view.

1. To add a new role, click **Add**.

   To change an existing role, click **Edit** for the firewall policy to be changed. The WebUI displays the **User Roles** tab on top.
2. **Role Name**. Enter the name for the role.
3. To add a policy for the new role, click **Add** in the Firewall Policies section. The WebUI expands the **Firewall Policies** section.

   Choose from existing configured policies, create a new policy based on existing policies, or create a new policy.
   a. If you elect to create a new policy, click on the radio button for **Create New Policy** and then click **Create**. The WebUI displays the **Policies** tab.
   b. In the Policies tab:

      Policy Name. Provide the policy name and select the IPv4 Session policy type from the drop-down list. The WebUI expands the **Policies** tab.
   c. In the drop-down lists, choose parameters such as source, destination, service in the same way as other firewall policy rules. For certain choices, the WebUI expands and adds drop-down lists.
   d. In the Action drop-down menu, select the **redirect to ESI group** option.
   e. In the Action drop-down menu, select the appropriate ESI group.
   f. Select the traffic direction. **Forward** refers to the direction of traffic from the (untrusted) client or user to the (trusted) server (such as the HTTP server or email server).

g. To add this rule to the policy, click **Add**.
h. Repeat the steps to configure additional rules.
i. Click **Done** to return to the **User Roles** tab. The WebUI returns to the **User Roles** tab.

4. Click **Apply** to apply the configuration changes.
5. Refer to Roles and Policies on page 364, for directions on how to apply a policy to a user role.

In the CLI

Use these commands to define the redirection filter for sending traffic to the ESI server and apply the firewall policy to a user role.

```
ip access-list session policy
   any any any redirect esi-group group direction both blacklist
   //For any incoming traffic, going to any destination,
   //redirect the traffic to servers in the specified ESI group.
   any any any permit
   //For everything else, allow the traffic to flow normally.

user-role role
   access-list {eth | mac | session}
   bandwidth-contract name
   captive-portal name
   dialer name
   pool {l2tp | pptp}
   reauthentication-interval minutes
   session-acl name
   vlan vlan_id
```

For example:

```
ip access-list session fortinet
   any any svc-http redirect esi-group fortinet direction both blacklist
   any any any permit

user-role guest
   access-list session fortinet
```

ESI Syslog Parser Domains and Rules

To configure the ESI syslog parser, navigate to the **Configuration > Advanced Services > External Services** view on the WebUI.The following sections describe how to manage syslog parser domains using the WebUI and CLI.

Managing Syslog Parser Domains in the WebUI

Click on the **Syslog Parser Domains** tab to display the Syslog Parser Domains view.

This view lists all the domains by domain name and server IP address, and includes a list of peer controllers (when peer controllers have been configured—as described in Understanding the ESI Syslog Parser on page 994).

Adding a new syslog parser domain

To add a new syslog parser domain:

1. Click **Add** in the **Syslog Parser Domains** section. The system displays the add domain view.
2. In the **Domain Name** text box, type the name of the domain to be added.
3. In the **Server IP Address** text box, type a valid IP address.

NOTE

You must ensure that you type a valid IP address, because the IP address you type is not automatically validated against the list of external servers that has been configured.

4. Click **Add**.
5. Click **Apply**.

Deleting an existing syslog parser domain

To delete an existing parser domain:

1. Identify the target parser domain in the list shown in the **Domain** section of the **Syslog Parser Domains** view.
2. Click **Delete** on the same row in the Actions column.

Editing an existing syslog parser domain

To change an existing syslog parser domain:

1. Identify the target parser domain in the list shown in the **Syslog Parser Domains** view .(see Managing Syslog Parser Domains in the WebUI on page 1000)
2. Click **Edit** on the same row in the **Actions** column.The system displays the edit domain view.

NOTE

You cannot modify the domain name when editing a parser domain.

3. To delete a server from the selected domain, highlight the server IP address and click **Delete**, then click **Apply** to commit the change.
4. To add a server or a peer controller to the selected domain, type the server IP address into the text box next to the **Add** button, click **Add**, then click **Apply** to commit the change, or click **Cancel** to discard the changes you made and exit the parser domain editing process.

When you make a change in the domain, you can click the **View Commands** link in the lower right corner of the window to see the CLI command that corresponds to the edit action you performed.

Managing Syslog Parser Domains in the CLI

Use these CLI commands to manage syslog parser domains.

Adding a new syslog parser domain

```
esi parser domain name
   peer peer-ip
   server ipaddr
```

Showing ESI syslog parser domain information

```
show esi parser domains
```

Deleting an existing syslog parser domain

```
no esi parser domain name
```

Editing an existing syslog parser domain

```
esi parser domain name
   no
   peer peer-ip
   server ipaddr
```

For example:

```
esi parser domain forti_domain
   server 30.0.0.1
   server 30.0.0.2
   server 30.0.0.3
   peer 20.0.0.1
```

Managing Syslog Parser Rules

The following sections describe how to manage syslog parser rules using the WebUI and CLI.

In the WebUI

Click on the **Syslog Parser Rules** tab to display the Syslog Parser Rules view.This view displays a table of rules with the following columns:

- Name— rule name
- Ena—where “y” indicates the rule is enabled and “n” indicates the rule is disabled (not enabled)
- Condition—Match condition (a regular expression)
- Match—Match type (IP address, MAC address, or user)
- User—Match pattern (a regular expression)
- Set—Set type (blacklist or role)
- Value—Set value (role name)
- Domain—Parser domain to which this rule is to be applied
- Actions—The actions that can be performed on each rule.

Adding a new parser rule

To add a new syslog parser rule:

1. Click **Add** in the **Syslog Parser Rules** view. The system displays the new rule view.
1. In the **Rule Name** text box, type the name of the rule you want to add.
2. Click the **Enable** checkbox to enable the rule.
3. In the **Condition Pattern** text box, type the regular expression to be used as the condition pattern.

   For example, “log_id=[0–9]{10}[ ]” to search for and match a 10-digit string preceded by “log_id=” and followed by one space.
4. In the drop-down **Match** list, use the drop-down menu to select the match type (ipaddr, mac, or user).
5. In the **Match Pattern** text box, type the regular expression to be used as the match pattern.

   For example, if you selected “mac” as the match type, type the regular expression to be used as the match pattern. You could use “mac[ ](.{17})” to search for and match a 17-character MAC address preceded by the word “mac” plus one space.
6. In the drop-down **Set** list, select the set type (blacklist or role).

   When you select **role** as the Set type, the system displays a second drop-down list. Click the list to display the possible choices and select the appropriate role value. Validation on the entered value will be based on the Set selection.
7. In the drop-down **Parser Group** list, select one of the configured parser domain names.

Deleting a syslog parser rule

To delete an existing syslog parser rule:

1. Identify the target parser rule in the list shown in the **Syslog Parser Rules** view.
2. Click **Delete** on the same row in the Actions column.

Editing an existing syslog parser rule

To change an existing syslog parser rule:

1. Identify the target parser rule in the list shown in the **Syslog Parser Rules** view.
2. Click **Edit** on the same row in the **Actions** column. The system displays the attributes for the selected rule



You cannot modify the rule name when editing a parser rule.

3. Change the other rule attributes as required:
   a. Click the **Enable** checkbox to enable the rule.
   b. In the **Condition Pattern** text box, type the regular expression to be used as the condition pattern.
   c. In the drop-down **Match** list, select the match type (ipaddr, mac, or user).
   d. In the **Match Pattern** text box, type the regular expression to be used as the match pattern.
   e. In the drop-down **Set** list, select the set type (blacklist or role).
   f. When you select **role** as the Set type, the system displays a second drop-down list. Click the list to display the possible choices and select the appropriate role value. Validation on the entered value will be based on the Set selection.
   g. In the drop-down **Parser Group** list, select one of the configured parser domain names.



At this point, you can test the rule you just edited by using the Test section of the edit rule view. You can also test rules outside the add or edit processes by using the rule test in the Syslog Parser Test view (accessed from the External Services page by clicking the Syslog Parser Test tab, described in Testing a Parser Rule on page 1003.

4. Click **Apply** to apply the configuration changes.

Testing a Parser Rule

You can test or validate enabled Syslog Parser rules against a sample syslog message, or against a syslog message file containing multiple syslog messages. Access the parser rules test from the **External Services** page by clicking the **Syslog Parser Test** tab, which displays the Syslog Parser Rule Test view.

To test against a sample syslog message:

   a. In the drop-down **Test Type** list, select **Syslog message** as the test source type.
   b. In the Message text box, type the syslog message text.
   c. Click **Test** to start the test.

      The test results are displayed in a box in the area below the Test button. The test results contain information about the matching rule and match pattern.

- To test against a syslog message file:
   a. In the drop-down **Test Type** list, select **Syslog file** as the test type.
   b. In the Filename text box, type the syslog file name.
   c. Click **Test** to start the test.

      The test results are displayed in a box in the area below the Test button. The test results contain information about the matching rule and match pattern.

In the CLI

Use these CLI commands to manage syslog parser rules.

Adding a new parser rule

```
esi parser rule rule-name
```

```
   condition expression
   domain name
   enable
   match {ipaddr expression | mac expression | user expression}
   position position
   set {blacklist | role role}
```

For example:

```
esi parser rule forti_virus
   condition "log_id=[0-9]{10}[ ]"
   match "src=(.*)[ ]"
   set blacklist
   enable
```

Showing ESI syslog parser rule information:

```
show esi parser rules
```

Deleting a syslog parser rule:

```
no esi parser rule rule-name
```

Editing an existing syslog parser rule

```
esi parser rule rule-name
   condition expression
   domain name
   enable
   match {ipaddr expression | mac expression | user expression}
   no
   position position
   set {blacklist | role role}
```

Testing a parser rule

```
esi parser rule rule-name
   test {file filename | msg message}
```

Monitoring Syslog Parser Statistics

The following sections describe how to monitor syslog parser statistics using the WebUI and CLI.

In the WebUI

You can monitor syslog parser statistics in the External Servers monitoring page, accessed by selecting **Monitoring > Switch > External Services Interface > Syslog Parser Statistics**.

The Syslog Parser Statistics view displays statistics such as the number of matches and number of users per rule, as well as the number of respective actions fired by the syslog parser.



The Syslog Parser Statistics view also displays the last refresh time stamp and includes a **Refresh Now** button, to allow the statistics information to be refreshed manually. There is no automatic refresh on this page.

In the CLI

```
show esi parser stats
```

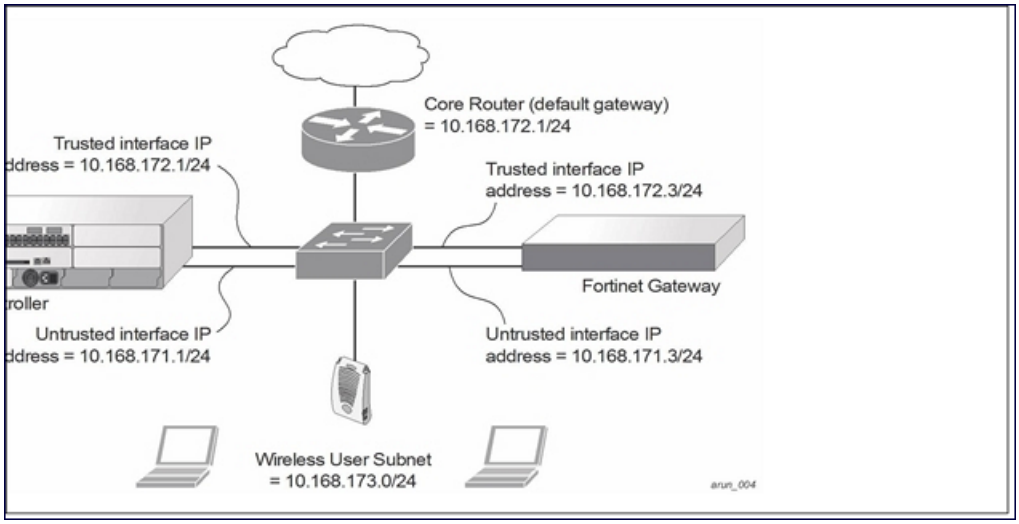
Sample Route-mode ESI Topology

This section introduces the configuration for a sample route-mode topology using the controller and Fortinet Anti-Virus gateways. In route mode, the trusted and untrusted interfaces between the controller and the Fortinet gateways are on different subnets. The following figure shows an example route-mode topology.

NOTE

ESI with Fortinet Anti-Virus gateways is supported only in route mode.

Figure 240 *Example Route-Mode Topology*



In the topology shown, the following configurations are entered on the controller and Fortinet gateway:

ESI server configuration on controller

- Trusted IP address = 10.168.172.3 (syslog source)
- Untrusted IP address = 10.168.171.3
- Mode = route

IP routing configuration on Fortinet gateway

- Default gateway (core router) = 10.168.172.1
- Static route for wireless user subnet (10.168.173.0/24) through the controller (10.168.171.2)

Configuring the Example Routed ESI Topology

This section describes how to implement the example routed ESI topology shown in . The description includes the relevant configuration—both the WebUI and the CLI configuration processes are described—required on the controller to integrate with a AVF server appliance.

The ESI configuration process will redirect all HTTP user traffic to the Fortinet server for examination, and any infected user will be blacklisted. The configuration process consists of these general tasks:

- Defining the ESI server.
- Defining the default ping health check method.
- Defining the ESI group.
- Defining the HTTP redirect filter for sending HTTP traffic to the ESI server.
- Applying the firewall policy to the guest role.
- Defining ESI parser domains and rules.

There are three configuration “phases” on the controller as a part of the solution.

- The first phase configures the ESI *ping health-check method*, *servers*, and *server groups*.The term *server* here refers to external AVF server devices.

- In the second phase of the configuration task, the user roles are configured with the redirection policies (session ACL definition) instructing the controller to redirect the different types of traffic to different server groups.
- In the final phase, the ESI parser domains and rules are configured.



The procedures shown in the following sections are based on the requirements in the example routed ESI topology. Your application might be broader or narrower than this example, but the same general operations apply.

Health-Check Method, Groups, and Servers

To configure the ESI health-check method, servers, and server groups, navigate to the **Configuration > Advanced Services > External Services** view on the WebUI.

Defining the Ping Health-Check Method

In the WebUI

To configure a health check profile:

1. Navigate to the **Configuration > Advanced Services > External Services** page on the WebUI.
2. Click **Add** in the **Health Check Configuration** section.
   To change an existing profile, click **Edit**.
3. Provide the following details:
   a. Enter enter the name **default for the Profile Name.**
   b. **Frequency (secs)**—Enter **5**.)
   c. **Timeout (secs)**—Indicates the number of seconds the controller waits for a response to its health check query before marking the health check as failed. Default: 2 seconds. (In this example, enter **3**.)
   d. **Retry count**—Is the number of failed health checks after which the controller marks the server as being down. Default: 2. (In this example, enter **3**.)
4. Click **Done** when you are finished.
5. Click **Apply**.

In the CLI

Use these CLI commands to configure a health-check method:

```
esi ping profile_name
   frequency seconds
   retry-count count
   timeout seconds
```

For example:

```
esi ping default
   frequency 5
   retry-count 3
   timeout 3
```

Defining the ESI Server

The following sections describe how to configure an ESI server using the WebUI and CLI.

In the WebUI

To configure an ESI server:

1. Navigate to the **Configuration > Advanced Services > External Services** page on the WebUI.

2. Click **Add** in the **External Servers** section.
3. Provide the following details:
   a. **Server Name**. (This example uses the name **forti_1**.)
   b. **Server Group**. Use the drop-down list to assign this server to a group from the existing configured groups. (This example uses **fortinet**.)
   c. **Server Mode**. Use the drop-down list to choose the mode (bridge, nat, or route) your topology requires. Refer to the description above to understand the differences between the modes. (This example uses **route** mode.)
   d. **Trusted IP Address**. Enter **10.168.172.3**.)
   e. **Untrusted IP Address**. Enter **10.168.171.3**.)
4. Click **Done** when you are finished.
5. Click **Apply** to apply the configuration changes.

In the CLI

Use these CLI commands to configure an ESI server and identify its associated attributes:

```
esi server server_identity
   dport destination_tcp/udp_port
   mode {bridge | nat | route}
   trusted-ip-addr ip-addr [health-check]
   trusted-port slot/port
   untrusted-ip-addr ip-addr [health-check]
   untrusted-port slot/port
```

For example:

```
esi server forti_1
   mode route
   trusted-ip-addr 10.168.172.3
   untrusted-ip-addr 10.168.171.3
```

Defining the ESI Server Group

The following sections describe how to configure an ESI server group using the WebUI and CLI.

In the WebUI

To configure an ESI server group on the controller:

1. Navigate to the **Configuration > Advanced Services > External Services** page.
2. Click **Add** in the **Server Groups** section.
3. Provide the following details:
   a. Enter a **Group Name.** Enter **fortinet**.)
   b. In the drop-down list, select **default** as the health check profile.
4. Click **Done** when you are finished.
5. Click **Apply** to apply the configuration changes.

In the CLI

Use these CLI commands to configure an ESI server group, identify its associated ping health-check method, and associate a server with this group:

```
esi group name
   ping profile_name
   server server_identity
```

For example:

```
esi group fortinet
   ping default
   server forti_1
```

Redirection Policies and User Role

The following sections describe how to configure the redirection policies and user role using the WebUI and CLI.

In the WebUI

To configure user roles to redirect the required traffic to the server(s), navigate to the **Configuration > Access Control > User Roles** view (see 1).

1. To add a new role, click **Add**.The WebUI displays the **Add Role** view.

   Role Name. Enter "guest" as the name for the role.
2. To add a policy for the new role, click **Add** in the Firewall Policies section. The WebUI expands the **Firewall Policies** section.

   Choose from existing configured policies, create a new policy based on existing policies, or create a new policy.

   a. If you elect to create a new policy, click on the radio button for **Create New Policy** and then click **Create**. The WebUI displays the **Policies** tab.

   b. In the Policies tab:

      Policy Name. Enter the policy name **fortinet** and the **IPv4 Session** policy type.) Click **Add** to proceed. The WebUI expands the **Policies** tab.

      In the drop-down lists, choose parameters such as source, destination, service in the same way as other firewall policy rules. This example uses **any** source, **any** destination, service type **svc-http (tcp 80)**. For certain choices, the WebUI expands and adds drop-down lists.

   c. In the Action drop-down menu, select the **redirect to ESI group** option.

      Select **fortinet** as the appropriate ESI group.

      The three steps above translate to "for any incoming HTTP traffic, going to any destination, redirect the traffic to servers in the ESI group named fortinet.")

      Select **both** as the traffic direction. **Forward** refers to the direction of traffic from the untrusted client or user to the trusted server, such as the HTTP server or email server.

      To add this rule to the policy, click **Add**.

   d. Repeat the steps to configure additional rules. This example adds a rule that specifies **any, any, any, permit**.

   e. Click **Done** to return to the **User Roles** tab.
3. Click **Apply** to apply the configuration changes.
4. Refer to Roles and Policies on page 364, for directions on how to apply a policy to a user role.

In the CLI

Use these commands to define the redirection filter for sending traffic to the ESI server and apply the firewall policy to a user role in the route-mode ESI topology example.

```
ip access-list session policy
   any any any redirect esi-group group direction both blacklist
   //For any incoming traffic, going to any destination,
   //redirect the traffic to servers in the specified ESI group.
   any any any permit
   //For everything else, allow the traffic to flow normally.
```

```
user-role role
   access-list {eth | mac | session}
   bandwidth-contract name
   captive-portal name
   dialer name
   pool {l2tp | pptp}
   reauthentication-interval minutes
   session-acl name
   vlan vlan_id
```

For example:

```
ip access-list session fortinet
   any any svc-http redirect esi-group fortinet direction both blacklist
   any any any permit
user-role guest
   access-list session fortinet
```

Syslog Parser Domain and Rules

The following sections describe how to configure the syslog parser domain and rules for the route-mode example using the WebUI and CLI.

Add a New Syslog Parser Domain in the WebUI

To add a new syslog parser domain for the routed example:

1. Click **Add** in the **Syslog Parser Domains** tab (**Advanced Services > External Services > Syslog Parser Domain**).

   The system displays the new domain view.
2. In the **Domain Name** text box, type the name of the domain to be added.
3. In the **Server (IP Address)** text box, type a valid IP address.



You must ensure that you type a valid IP address, because the IP address you type is not automatically validated against the list of external servers that has been configured.

4. Click **<< Add**.
5. Click **Apply**.

Adding a New Parser Rule in the WebUI

To add a new syslog parser rule for the route-mode example:

1. Click **Add** in the **Syslog Parser Rules** tab (**Advanced Services > External Services > Syslog Parser Rule**).

   The system displays the new rule view.
2. In the **Rule Name** text box, type the name of the rule to be added (in this example, “forti_virus”).
3. Click the **Enable** checkbox to enable the rule.
4. In the **Condition Pattern** text box, type the regular expression to be used as the condition pattern. (In this example, the expression “log_id=[0–9]{10}[ ]” searches for and matches a 10-digit string preceded by “log_id=” and followed by one space.)
5. In the drop-down **Match** list, use the drop-down menu to select the match type (in this example, ipaddr).
6. In the **Match Pattern** text box, type the regular expression to be used as the match pattern (in this example, “src=(.*)[ ]”).
7. In the drop-down **Set** list, select the set type (in this example, blacklist).

8. In the drop-down **Parser Group** list, select one of the configured parser domain names (in this example, “forti_domain”).
9. Click **Apply**.

In the CLI

Use these CLI commands to define a syslog parser domain and the rule to be applied in the route-mode example shown in Figure 240

```
esi parser domain name
   peer peer-ip
   server ipaddr

esi parser rule rule-name
   condition expression
   domain name
   enable
   match {ipaddr expression | mac expression | user expression }
   position position
   set {blacklist | role role}
```

For example:

```
esi parser domain forti_domain
   server 10.168.172.3

esi parser rule forti_virus
   condition “log_id=[0-9]{10}[ ]”
   match ipaddr “src=(.*)[ ]”
   set blacklist
   enable
```

Sample NAT-mode ESI Topology

This section describes the configuration for a sample NAT-mode topology using the controller and three external captive-portal servers. NAT mode uses a trusted interface for each external captive-portal server and a different destination port to redirect a packet to a port other than the original destination port in the packet. An example topology is shown below in Figure 242.

Figure 241 *Example NAT-Mode Topology*

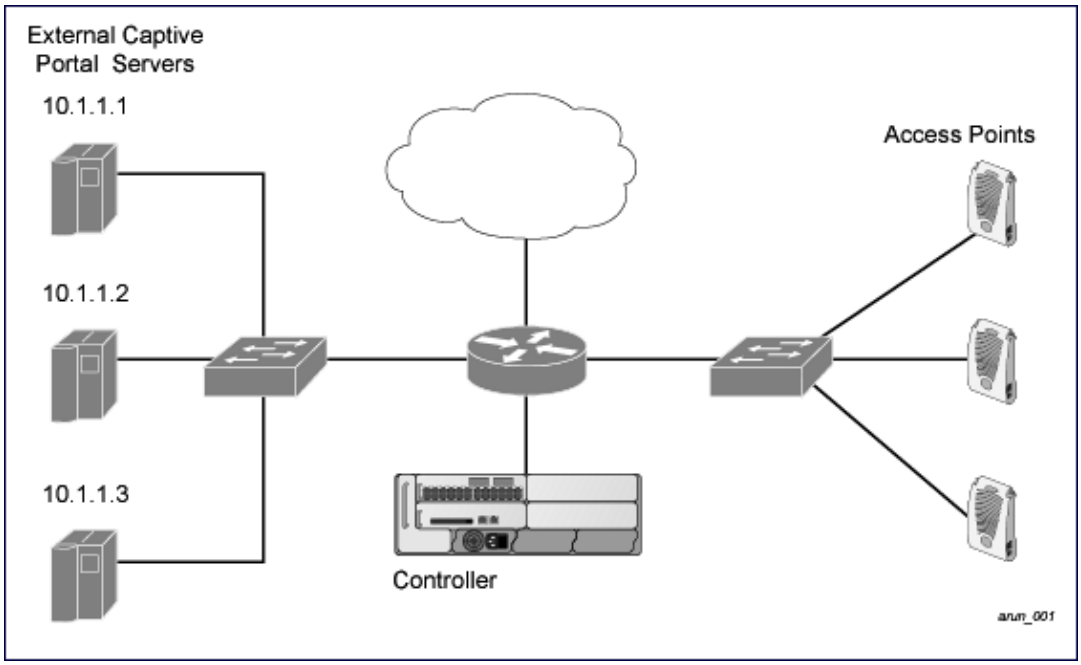
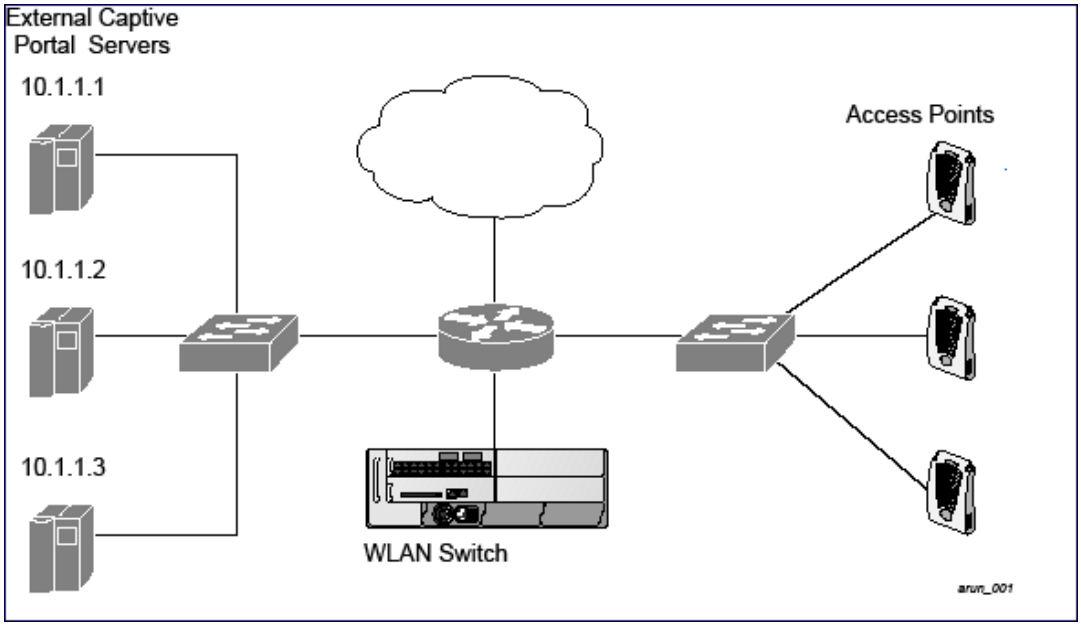


Figure 242



In this example, all HTTP traffic received by the controller is redirected to the external captive portal server group and load-balanced across the captive portal servers. All wireless client traffic with destination port 80 is redirected to the captive portal server group, with the new destination port 8080.

NOTE

The external servers do not necessarily have to be on the subnet as the controller. The policy that redirects traffic to the external servers for load balancing is routed to the external servers if they are on a different subnet.

In the topology shown, the following configurations are entered on the controller and external captive-portal servers:

ESI server configuration on the controller

- External captive-portal server 1:
  - Name = external_cp1
  - Mode = NAT
  - Trusted IP address = 10.1.1.1
  - Alternate destination port = 8080
- External captive-portal server 2:
  - Name = external_cp2
  - Mode = NAT
  - Trusted IP address = 10.1.1.2
- External captive-portal server 3:
  - Name = external_cp3
  - Mode = NAT
  - Trusted IP address = 10.1.1.3
- Health-check ping:
  - Name = externalcp_ping
  - Frequency = 30 seconds
  - Retry-count = 2 attempts
  - Timeout = 2 seconds (2 seconds is the default)
- ESI group = external_cps
- Session access control list (ACL)
  - Name = cp_redirect_acl
  - Session policy = user any svc-http redirect esi-group external_cps direction both

Configuring the Example NAT-mode ESI Topology

This section describes how to implement the example NAT-mode ESI topology shown in using both the WebUI, then the CLI.

The configuration process consists of these general tasks:

- Configuring captive portal (see the “Configuring Captive Portal” chapter).
- Configuring the health-check ping method.
- Configuring the ESI servers.
- Configuring the ESI group.
- Defining the redirect filter for sending traffic to the ESI server.

Configuring the NAT-mode ESI Example in the WebUI

Navigate to the **Configuration > Advanced Services > External Services** view on the WebUI (see).

In the WebUI

1. Click **Add** in the **Health-Check Configuration** section **External Services** view on the WebUI.
2. Provide the following details:
   a. **Profile Name**. This example uses **externalcp_ping**.

b. **Frequency** seconds. This example uses **30**.
c. **Retry Count**. This example uses **3**.



If you do not specify a value for a parameter, the WebUI assumes the default value. In this example, the desired timeout value is two seconds; therefore, not specifying the timeout causes the WebUI to use the default value of two seconds.

3. Click **Done** when you are finished.



To apply the configuration (changes), you must click **Apply** in the **External Services** view on the WebUI. In this example, you can wait until you finish configuring the servers and groups, or you can apply after each configuration portion.

Configuring the ESI Group in the WebUI

1. Click **Add** in the **Server Groups** section **External Services** view on the WebUI.
2. Provide the following details:
   a. **Group Name**. This example uses **external_cps**.
   b. **Health-Check Profile**. Select the health-check ping from the drop-down list. This example uses **externalcp_ping**.
3. Click **Done** when you are finished.



To apply the configuration (changes), you must click **Apply** in the **External Services** view on the WebUI. In this example, you can wait until you finish configuring the servers and groups, or you can apply after each configuration portion.

Configure the ESI Servers in the WebUI

1. Click **Add** in the **External Servers** section.
2. Provide the following details:
   a. **Server Name**.
   b. **Server Group**. Use the drop-down list to assign this server to a group from the existing configured groups.
   c. **Server Mode**. Use the drop-down list to choose NAT mode.)
   d. **Trusted IP Address**. For nat mode, enter the IP address of the trusted interface on the external captive portal server.
   e. **NAT Destination Port**. Enter the port number (to redirect a packet to a port other than the original destination port in the packet).
3. Click **Done** when you are finished.
4. Repeat Step 1 through Step 3 for the remaining external captive portal servers.
5. Click **Apply** to apply the configuration changes.

Configuring the Redirection Filter in the WebUI

To redirect the required traffic to the server(s) using the WebUI, navigate to the **Configuration > Access Control > User Roles** view on the WebUI (see 1).

1. Click the **Policies** tab.
2. Click **Add** in the **Policies** section of the **Policies** view on the WebUI.
3. Provide the following details:
   a. **Policy Name**. (This example uses **cp_redirect_acl**.)

b. **Policy Type**. Select **IPv4 Session** from the drop-down list.

4. Click **Add** in the **Rules** section of the **Policies** view.
   a. **Source**. Select **user** from the drop-down list.
   b. **Destination**. Accept **any**.
   c. **Service**. Select **service** from the drop-down list; select **svc-http (tcp 80)** from the secondary drop-down list.
   d. **Action**. Select **redirect to ESI group** from the drop-down list; select **external_cps** from the secondary drop-down list; click **<--** to add that group.
   e. Click **Add**.
5. Click **Done** when you are finished.
6. Click **Apply** to apply the configuration changes.

Configuring the Example NAT-mode Topology in the CLI

The CLI configuration process consists of these general tasks:

- Configuring captive portal (see Captive Portal Authentication on page 299).
- Configuring the health-check ping method.
- Configuring the ESI servers.
- Configuring the ESI group.
- Defining the redirect filter for sending traffic to the ESI server.

Configuring a Health-Check Ping

The health-check ping will be associated with an ESI group, along with servers, so that controller will send ICMP echo requests to each server in the group and mark the server down if the controller does not hear from the server. The health-check parameters used in this example are:

- Frequency—30 seconds. (The default is 5 seconds.)
- Retry-count—3. (The default is 2.)
- Timeout—2 seconds. (The default is 2 seconds.)

Use these CLI commands to configure a health-check ping method:

```
esi ping profile_name
   frequency seconds
   retry-count count
   timeout seconds
```

Configuring ESI Servers

Here are the ESI server CLI configuration tasks:

- Configure server mode to be NAT.
- Configure the trusted IP address (the server IP address to which packets should be redirected).
- To redirect to a different port than the original destination port in the packet, configure an alternate destination port.

Use these CLI commands to configure an ESI server and identify its associated attributes:

```
esi server server_identity
   dport destination_tcp/udp_port
   mode {bridge | nat | route}
   trusted-ip-addr ip-addr [health-check]
```

Configure an ESI Group, Add the Health-Check Ping and ESI Servers

Use these CLI commands to configure an ESI server group, identify its associated ping health-check method, and associate a server with this group:

```
esi group name
   ping profile_name
   server server_identity
```

Using the ESI Group in a Session Access Control List

Use these CLI commands to define the redirection filter for sending traffic to the ESI server.

```
ip access-list session policy
   user any svc-http redirect esi-group group direction both
```

CLI Configuration Example 1

```
esi ping externalcp_ping
   frequency 30
   retry-count 3

esi server external_cp1
   dport 8080
   mode nat
   trusted-ip-addr 10.1.1.1

esi server external_cp2
   dport 8080
   mode nat
   trusted-ip-addr 10.1.1.2

esi server external_cp3
   dport 8080
   mode nat
   trusted-ip-addr 10.1.1.3

esi group external_cps
   ping externalcp_ping
   server external_cp1
   server external_cp2
   server external_cp3

ip access-list session cp_redirect_acl
   user any svc-http redirect esi-group external_cps direction both
```

CLI Configuration Example 2

```
esi server https-proxy1
   dport 44300
   mode nat
   trusted-ip-addr 1.2.3.4

esi server https-proxy2
   dport 44300
   mode nat
   trusted-ip-addr 1.2.3.5

esi group https-proxies
   ping default
   server https-proxy1
   server https-proxy2

ip access-list session https-proxy
```

```
user any svc-https redirect esi-group https-proxies direction both
any any permit
```

Understanding Basic Regular Expression (BRE) Syntax

The ESI syslog parser supports regular expressions created using the Basic Regular Expression (BRE) syntax described in this section. BRE syntax consists of instructions—character-matching operators (described in Table 220), repetition operators (described in Table 221), or expression anchors (described in Table 222)—used to defined the search or match target.

This section contains the following topics:

- "Character-Matching Operators" on page 512
- "Regular Expression Repetition Operators" on page 513
- "Regular Expression Anchors" on page 513
- "References" on page 514

Character-Matching Operators

Character-matching operators define what the search will match.

Table 220: *Character-matching operators in regular expressions*

Operator	Description	Sample	Result
.	Match any one character.	grep .ord sample.txt	Matches *ford*, *lord*, *2ord*, etc. in the file sample.txt.
[ ]	Match any one character listed between the brackets	grep [cng]ord sample.txt	Matches only *cord*, *nord*, and *gord*
[^]	Match any one character not listed between the brackets	grep [^cn]ord sample.txt	Matches *lord*, *2ord*, etc., but not *cord* or *nord*
		grep [a-zA-Z]ord sample.txt	Matches *aord*, *bord*, *Aord*, *Bord*, etc.
		grep [^0-9]ord sample.txt	Matches *Aord*, *aord*, etc., but not *2ord*, etc.

Regular Expression Repetition Operators

Repetition operators are *quantifiers* that describe how many times to search for a specified string. Use them in conjunction with the character-matching operators in Table 221 to search for multiple characters.

Table 221: *Regular expression repetition operators*

Operator	Description	Sample	Result
?	Match any character one time if it exists	egrep "?erd" sample text	Matches *berd*, *herd*, etc., *erd*
*	Match declared element multiple times if it exists	egrep "n.*rd" sample.txt	Matches *nerd*, *nrd*, *neard*, etc.
+	Match declared element one or more times	egrep "[n]+erd" sample.txt	Matches *nerd*, *nnerd*, etc., but not *erd*
{n}	Match declared element exactly *n* times	egrep "[a-z]{2}erd" sample.txt	Matches *cherd*, *blerd*, etc., but not *nerd*, *erd*, *buzzerd*, etc.
{n,}	Match declared element at least *n* times	egrep ".{2,}erd" sample.txt	Matches *cherd* and *buzzerd*, but not *nerd*
{n,N}	Match declared element at least *n* times, but not more than *N* times	egrep "n[e]{1,2}rd" sample.txt	Matches *nerd* and *neerd*

Regular Expression Anchors

Anchors describe where to match the pattern, and are a handy tool for searching for common string combinations. Some of the anchor examples use the vi line editor command *:s*, which stands for *substitute*. That command uses the syntax: `s/pattern_to_match/pattern_to_substitute`.

Table 222: *Regular expression anchors*

Operator	Description	Sample	Result
^	Match at the beginning of a line	s/^/blah /	Inserts "blah" at the beginning of the line
$	Match at the end of a line	s/$/ blah/	Inserts " blah" at the end of the line
\<	Match at the beginning of a word	s/\</blah/	Inserts "blah" at the beginning of the word
		egrep "\<blah" sample.txt	Matches *blahfield*, etc.
\>	Match at the end of a word	s/\>/blah/	Inserts "blah" at the end of the word
		egrep "\>blah" sample.txt	Matches *soupblah*, etc.
\b	Match at the beginning or end of a word	egrep "\bblah" sample.txt	Matches *blahcake* and *countblah*
\B	Match in the middle of a word	egrep "\Bblah" sample.txt	Matches *sublahper*, etc.

References

This implementation is based, in part, on the following resources:

- Lonvick, C., “The BSD syslog Protocol”, RFC 3164, August 2001
- Regular expression (regex) reference: en.wikipedia.org/wiki/Regular_expression
- Regex syntax summary:greenend.org.uk/rjk/2002/06/regexp.html
- Basic regular expression (BRE) syntax: builder.com.com/5100-6372-1050915.html

This chapter introduces the ArubaOS XML API interface and briefly discusses how you can use the simple API calls to perform external user management tasks. A sample code listing at the end of the chapter to help you get started with using the XML API.

Topics in this chapter include:

- Overview on page 1019
- Working with the ArubaOS XML API Works on page 1019
- Creating an XML Request on page 1019
- XML Response on page 1021
- Sample Code on page 1029

Overview

ArubaOS allows you to set up customized external captive portal user management using its native XML API interface. The XML API interface allows you to create and execute user management operations on behalf of the clients or users. You can use the XML API interface to add, delete, authenticate, or query a user or a client.

Before you Begin

- Enable the External Services Interface software module. This is available in the PEFNG license.
- Ensure that you have connectivity between your captive portal server and the controllers via HTTP or HTTPS.

Working with the ArubaOS XML API Works

The typical interaction between your external server and the controller happens over HTTPS post commands. A typical communication process using the XML API interface happens as follows:

1. An API command is issued from your server in XML format to the controller. The XML message or request can be composed using a language of your choice using the format described in the Creating an XML Request on page 1019. Sample code in *C* gives a simple example. See the Sample Code on page 1029.
2. The controller processes the XML request and sends the response to the authentication server in the XML format. The XML request is sent using HTTPS post. The common format of the HTTPS post is `https://<controller-ip>/auth/command.xml`. See Creating an XML Request on page 1019 for more information.
3. You can use the response and take appropriate action that suit your requirements. The response from the controller is returned using predefined formats. See the XML Response on page 1021 for more information.

Creating an XML Request

You can create XML requests to add, delete, authenticate, blacklist, or query a user. This section provides XML request formats that you can use for each authentication tasks.

NOTE

The XML API functionalities such as addition, deletion, role change, querying, authentication, and blacklisting has been extended to support IPv6 users in addition to IPv4 users. The XML API server is configured using only the IPv4 address.

Adding a User

This XML requests uses the `user_add` command to create a new user entry in the controllers user table. If the user entry is already present in the user table, the command will modify the entry with the values defined in the XML request.

```
xml=<aruba command="user_add">
  <ipaddr>IP-address_of_the_user</ipaddr>
  <macaddr>MAC-address_of_the_user</macaddr>
  <name>User_Name</name>
  <role>Role_Name<role>
  <session_timeout>Session_timeout</session_timeout>
  <key>Shared_Key</key>
  <authentication>MD5|SHA-1|cleartext</authentication>
  <version>1.0</version>
</aruba>
```

The following options are mandatory when you execute the `user_add` command:

- IP Address
- Version

Deleting a User

This XML requests uses the `user_delete` command to delete an existing user from the controllers user table. If the user entry contains multiple attributes these must be specified in the XML request

```
xml=<aruba command="user_delete">
  <ipaddr>IP-address_of_the_user</ipaddr>
  <macaddr>MAC-address_of_the_user</macaddr>
  <name>User_Name</name>
  <key>Shared_Key</key>
  <authentication>MD5|SHA-1|cleartext</authentication>
  <version>1.0</version>
</aruba>
```

The following options are mandatory when you execute the `user_add` command:

- IP Address
- Version

Authenticating a User

This XML requests uses the `user_authenticate` command to authenticate and derive a new for the user.

```
xml=<aruba command="user_authenticate">
  <ipaddr>IP-address_of_the_user</ipaddr>
  <macaddr>MAC-address_of_the_user</macaddr>
  <name>User_Name</name>
  <password>Password_for_the_user</password>
  <key>Shared_Key</key>
  <authentication>MD5|SHA-1|cleartext</authentication>
  <version>1.0</version>
</aruba>
```

The following options are mandatory when you execute the `user_authenticate` command:

- IP Address
- Version
- Name
- Password

Blacklisting a User

This XML requests uses the `user_blacklist` command to blacklist a user from connecting to your network.

```
xml=<aruba command="user_blacklist">
  <ipaddr>IP-address_of_the_user</ipaddr>
  <macaddr>MAC-address_of_the_user</macaddr>
  <name>User_Name</name>
  <key>Shared_Key</key>
  <authentication>MD5|SHA-1|cleartext</authentication>
  <version>1.0</version>
</aruba>
```

The following options are mandatory when you execute the `user_blacklist` command:

- IP Address
- Version

Querying for User Status

This XML requests uses the `user_query` command to get the status and details of a user connected to your network.

```
xml=<aruba command="user_query">
  <ipaddr>IP-address_of_the_user</ipaddr>
  <macaddr>MAC-address_of_the_user</macaddr>
  <name>User_Name</name>
  <key>Shared_Key</key>
  <authentication>MD5|SHA-1|cleartext</authentication>
  <version>1.0</version>
</aruba>
```

The following options are mandatory when you execute the `user_blacklist` command:

- IP Address
- Version

XML Response

For every successful XML request the controller will return the processed information as an XML response. There are two types of responses: Default response and Query response.

Default Response Format

The format of a default XML response from the controller is:

```
<aruba>
  <result>Error | Ok</result>
  <code>response_code</code>
  <reason>response_message</reason>
</aruba>
```

In which,

- Result specifies if the XML result was successful or failure. If the request was successful, the result tag will contain the **`Ok`** string. If the request was a failure, the result tag will contain the **`Error`** string.
- Code is an integer number that represents the error in the request. This tag is populated only if there is an error in the request.
- Reason is message that contain descriptive information about error.

Response Codes

The following response codes are returned if the XML request return an the Error string.

Table 223: *XML Response Codes*

Code	Reason message	Description
1	`unknown user` The user specified in the XML request does not exist or is incorrect.	Returned by the `user_authenticate`, `user_delete`, `user_blacklist`, and `user_query` commands.
2	**`unknown role`** The specified role in the XML request does not exist in the controller.	Returned by the `user_add` command.
3	`unknown external agent`	Returned by all commands.
4	`authentication failed` The username and the key does not match.	Returned by commands that contain the shared_key in XML request.
5	`invalid command` The XML request contains a command not supported by ArubaOS XML API interface.	—
6	`invalid message authentication method` The authentication method specified in the XML request is not supported by the ArubaOS XML API interface.	Returned by commands that contain the authentication method in the XML request.
7	`invalid message digest`	Returned by commands that contain the shared_key in the XML request.
8	`missing message authentication` The authentication method is not specified in the XML request.	Returned by all commands that require the authentication method in the XML request.
9	**`missing or invalid version number`** The XML request does not contain the version number or the version number is incorrect.	Returned by all commands.
10	`internal error`	
11	`client not authorized` The shared key in the XML request does not match or the XML API server is not defined in the appropriate AAA profile.	Returned by all commands that require shared key to be specified in the XML request.
12	`Cant use VLAN IP`	—

Code	Reason message	Description
13	`Invalid IP` The XML request contains invalid IP address of the user or client.	Returned by all commands that required IP address to be specified in the XML request.
14	`Cant use Switch IP` The XML request contains the controllers IP address instead of the client IP address.	Returned by all commands that required IP address to be specified in the XML request.
15	`missing MAC address` The XML request does not contain the MAC address of the user or client.	Returned by all commands that required MAC address to be specified in the XML request.
16	`Unsupported command for this user`	Returned when the requested operation is invalid for the specified user.
17	`Socket failed or timed out waiting for operation to complete`	Returned when the status of the requested operation is unavailable; usually signifies a socket communication failure or timeout.

Query Command Response Format

The response of the XML request with the user_query command contains detailed information about the status of the user or client. The format of the response of a query command is:

```
<aruba>
  <result>Result</result>
  <code>Code</code>
  <reason>Reason</reason>
  <role>Role</role>
  <type>Type</type>
  <auth_status>Auth_status</auth_status>
  <auth_server>Auth_server</auth_server>
  <auth_method>Auth_method</auth_method>
  <location>Location</location>
  <age>Age</age>
  <essid>Essid</essid>
  <bssid>Bssid</bssid>
  <phy_type>Phy_type</phytype>
  <vlan>Vlan</vlan>
</aruba>
```

In which, the result, code and reason values are similar to the default response. The following responses, however, are returned only in the result code returns the `OK` string.

Table 224: *Query Response Code*

Response Code	Description
Role	Displays the current role of the authenticated user
Type	Displays is the user or client is **`wired`** or **`wireless`**.
Auth_status	Displays the authentication status of the user or client. Available values are:

Response Code	Description
	authenticated or **unauthenticated**.
Auth_server	Displays the name of the authentication server used for authenticating the user. This information is available only if the user is authenticated by the controller.
Auth_method	Displays the authentication mechanism used to authenticate the user. This information is available only if the user is authenticated by the controller.
Location	Displays the current location of the user / clients. For wireless clients, the location is displayed in the B.F.L format. For wired clients, the location is displayed in the slot/port format.
Age	Displays the age of user in the controller. The age is displayed in DD:HH:MM format (Day:Hours:Minutes).
ESSID	Displays the ESSID to which the user is associated.
BSSID	Displays the BSSID of the AP to which the user is associated.
Phy Type	Displays the physical connection type. One of a, b, or g.
Vlan	Displays the VLAN ID of the user.

Using the XML API Server

To use the XML API:

1. Configure an external XML API server
2. Associate the XML API server to an appropriate AAA profile
3. Configure a user role to direct un-authenticated users to the external captive portal server
4. Configure Captive Portal profile and associate that to an initial role (example `logon`)
5. Create an XML request with the appropriate API call
6. Process XML response appropriately



The default logon role of a client or user must have captive-portal enabled.

Configuring the XML API Server

Configure an external XML API server in your AAA infrastructure. In this example, **10.11.12.13** is your server. The XML API interface on the controller will receive requests from this server.

- Define the XML API server and specify the key for verifying requests from your server

```
(host) (config) #aaa xml-api server 10.11.12.13
(host) (XML API Server "10.11.12.13") #key $abcd$1234$
```

- Verify the XML API server configuration

```
(host) (config) #show aaa xml-api server
XML API Server List
-------------------
Name        References  Profile Status
----        ----------  --------------
```

```
10.11.12.13 1
is incremented for each usage.

Total:1
```

Associating the XML API Server to a AAA profile

After you define the XML API server profile associate it to the appropriate AAA profile. If the XML API server is not correctly configured in the appropriate profile, the controller will respond with the `client not authorized` error message. You can add XML API server references to the following AAA profile depending on your requirement:

- For wireless users—Associate the XML API server to the AAA profile of the virtual AP profile.

```
(host) (config) #aaa profile wirelessusers
(host) (AAA Profile "wirelessusers") #xml-api-server 10.11.12.13
(host) (XML API Server "10.11.12.13") #key Dell123
(host) (config) #show aaa profile wirelessusers

AAA Profile "wirelessusers"
---------------------------
Parameter                          Value
---------                          -----
Initial role                       logon
MAC Authentication Profile         N/A
MAC Authentication Default Role    guest
MAC Authentication Server Group    default
802.1X Authentication Profile      N/A
802.1X Authentication Default Role guest
802.1X Authentication Server Group N/A
RADIUS Accounting Server Group     N/A
XML API server                     10.11.12.13
RFC 3576 server                    N/A
User derivation rules              N/A
Wired to Wireless Roaming          Enabled
SIP authentication role            N/A

(host) (config) #wlan virtual-ap wireless-vap
(host) (Virtual AP profile "wireless-vap") #aaa-profile wirelessusers
(host) (config) #show wlan virtual-ap wireless-vap

Virtual AP profile "wireless-vap"
---------------------------------
Parameter                                       Value
---------                                       -----
Virtual AP enable                               Enabled
Allowed band                                    all
AAA Profile                                     wirelessusers
802.11K Profile                                 default
SSID Profile                                    default
VLAN                                            N/A
Forward mode                                    tunnel
Deny time range                                 N/A
Mobile IP                                       Enabled
HA Discovery on-association                     Disabled
DoS Prevention                                  Disabled
Station Blacklisting                            Enabled
Blacklist Time                                  3600 sec
Dynamic Multicast Optimization (DMO)            Disabled
Dynamic Multicast Optimization (DMO) Threshold  6
Authentication Failure Blacklist Time           3600 sec
Multi Association                               Disabled
```

```
Strict Compliance                                Disabled
VLAN Mobility                                    Disabled
Remote-AP Operation                              standard
Drop Broadcast and Multicast                     Disabled
Convert Broadcast ARP requests to unicast        Disabled
Band Steering                                    Disabled
WMM Traffic Management Profile                   N/A
```

- For wired users—Associate the XML API server to the AAA profile of the appropriate wired profile.

```
(host) (config) #aaa profile wiredusers
(host) (AAA Profile "wiredusers") #xml-api-server 10.11.12.13
(host) (AAA Profile "wiredusers") #!
(host) (config) #aaa authentication wired
(host) (Wired Authentication Profile) #profile wiredusers
(host) (Wired Authentication Profile) #show aaa authentication wired

Wired Authentication Profile
----------------------------
Parameter    Value
---------    -----
AAA Profile  wiredusers
```

- Unknown wired users—Associate the XML API server to the `default-xml-api` AAA profile.



The `default-xml-api` AAA profile is used only to add or authenticate new users.

The following example illustrates using the `default-xml-api` AAA profile.

```
(host) (config) #aaa profile default-xml-api
(host) (AAA Profile "default-xml-api") #xml-api-server 10.11.12.13
(host) (config) #show aaa profile default-xml-api
AAA Profile "default-xml-api" (Predefined (changed))
----------------------------------------------------
Parameter                          Value
---------                          -----
Initial role                       logon
MAC Authentication Profile         N/A
MAC Authentication Default Role    guest
MAC Authentication Server Group    default
802.1X Authentication Profile      N/A
802.1X Authentication Default Role guest
802.1X Authentication Server Group N/A
RADIUS Accounting Server Group     N/A
XML API server                     10.11.12.13
RFC 3576 server                    N/A
User derivation rules              N/A
Wired to Wireless Roaming          Enabled
SIP authentication role            N/A
```

Your controller is now ready to receive API calls from your XML API server.

Set up Captive Portal profile

Set up a Captive Portal profile with a login page that will redirect users to the external Captive Portal server.

```
(host) (config-role) #aaa authentication captive-portal captive-portal-auth
(host) (Captive Portal Authentication Profile "captive-portal-auth") #default-role
authenticated
(host) (Captive Portal Authentication Profile "captive-portal-auth") #login-page
https://10.11.12.13/cgi-bin/login.pl
```

```
(host) (Captive Portal Authentication Profile "captive-portal-auth") #switch-in-redirection-
url
```

Associating the Captive Portal Profile to an Initial Role

```
(host) (Captive Portal Authentication Profile "captive-portal-auth") #user-role logon

(host) (config-role) #captive-portal captive-portal-auth
(host)(config-role) #session-acl captiveportal
```

You can either create a new ACL or append specific rules to an exisiting ACLs. To create session ACL for the logon role do the following:

```
(host) (config-role) #netdestination xCP #an alias for the external Captive Portal server
(host) (config-dest) #host 10.11.12.13 #IP address of the external Captive Portal server
(host) (config-dest) #ip access-list session captiveportal #append or add rules to session ACL
(host) (config-sess-captiveportal)#user alias xCP svc-https permit
(host) (config-sess-captiveportal)#user alias xCP svc-http permit
```

Creating an XML API Request

You can now create an XML request with an appropriate authentication command and send it to the controller via HTTPS post. The format of the URL to send the XML request is:

```
https://<controller-ip/auth/command.xml
```

- `controller-ip` is the IP address of the controller that will receive the authentication request
- `command.xml` is the XML request that contains the details of authentication.

The format of the XML API request is:

```
xml=<aruba command="<authentication_command>">
  <options>Value</options>
  ...
  <options>Value</options>
</aruba>
```

You can specify any of the following commands in the XML request:

Table 225: *XML API Authentication Command*

Authentication Command	Description
`user_add`	This command adds the user to the controllers user table.
`user_delete`	This command deletes the user from the controller
`user_authenticate`	This command will authentication the user based on the authentication rules defined in the controllers configuration.
`user_blacklist`	This command will block a user from connection to your network.
`user_query`	This command will display the current status of the user connected to your network.

The authentication command requires certain mandatory options to successfully execute the authentication tasks. The list of all available options are:

Table 226: *Authentication Command Options*

Options	Description	Range / Defaults
`ipaddr`	IP address of the user in A.B.C.D format.	—
`macaddr`	MAC address of the user aa:bb:cc:dd:ee:ff format.	Enter MAC address with colon.
`user`	Name of the user.	64 character string
`role`	Role name assigned after authenticating.	64 character string
`password`	The password of the user used for authentication.	—
`session_timeout`	Session time-out in seconds. User will be disconnected after this time.	—
`authentication`	Authentication method used to authenticate the message and the sender. You can use any of MD5, SHA-1 or clear text methods of authentication. This option is ignored if shared secret is not configured. It is, however, mandatory if it is configured.	—
`key`	This is the encoded SHA1/MD5 hash of shared secret or plaintext shared secret. This option is ignored if shared secret is not configured on the switch. The actual MD5/SHA-1 hash is 16/20 bytes and consists of binary data. It must be encoded as an ASCII based HEX string before sending. It must be present when the controller is configured with an xml-api key for the server. Encoded hash length is 32/40 bytes for MD5/SHA-1.	
`version`	The version of the XML API interface available in the controller. This field is mandatory is all requests.	Current version 1.0

Monitoring External Captive Portal Usage Statistics

To check the external captive portal authentication statistics use the `show aaa xml-api statistics` command. This command displays the number of times an authentication command was executed per client. The command also displays the number of times an authentication event occurred and the number of new authentication events that occurred since the last status check.

```
(host) # show aaa xml-api statistics
ECP Statistics
--------------
Statistics                          10.10.10.249
----------                          ----------
user_authenticate                   1 (0)
user_add                            1 (0)
user_delete                         1 (0)
user_blacklist                      2 (0)
unknown user                        2 (0)
unknown role                        0 (0)
unknown external agent              0 (0)
authentication failed               0 (0)
invalid command                     0 (0)
invalid message authentication method  0 (0)
```

```
invalid message digest                 0 (0)

Packets received from unknown clients : 0 (0)
Packets received with unknown request : 0 (0)
Requests Received/Success/Failed      : 5/3/2 (0/0/0)
```

Sample Code

This section lists a sample code that will help you get started in using the ArubaOS XML API interface. These codes have been tested in a controlled environment. We recommend that you test this code in a non-production environment before using it for actual user management tasks.

Using XML API in C Language

The example script is written in the *C* language. The example script (`auth.c`) sends an authentication request from your authentication server to the controller.



This is an example code and is provided for illustration purposes. If you plan to use this code in your environment, ensure that the code meets your IT guidelines. Also create an error free executable to successfully execute the script.

Figure 243 *Authentication Script Listing*

```
##### auth.c listing
##### Authentication Script Example -- Start --
#include <stdio.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <getopt.h>

char *command, *ipaddr, *macaddr;
char *name, *password, *role;
char *tout, *secret;
char *auth, *key, enchashbuf[41];
unsigned char hashbuf[20];
char *version;

char post[4096], cmdbuf[512], encbuf[1024];

#define DEBUG
#ifdef DEBUG
#define debug(x...)     fprintf(stderr, x)
#else
#define debug(x...)
#endif

extern int cgi_escape_url(char *t, int tl, char *s, int sl, int b_newline);
static void encode_message_digest (unsigned char *md, int mdlen, char *output);

static void usage (void)
{
        fprintf(stderr, "Usage: ecp [options] <switch> <command> [<secret>]\n");

        fprintf(stderr, "  \n");
        fprintf(stderr, "  <switch>     Switch IP address.\n");
        fprintf(stderr, "  <command>    One of add, del, or authenticate.\n");
        fprintf(stderr, "  <secret>     Shared secret.\n");
        fprintf(stderr, "  \n");

        fprintf(stderr, "  -i ipaddr    User IP address in A.B.C.D format.\n");
```

```
        fprintf(stderr, "  -m macaddr   User MAC address in aa:bb:cc:dd:ee:ff format.\n");
        fprintf(stderr, "  -n name      User name.\n");
        fprintf(stderr, "  -p passwd    User password.\n");
        fprintf(stderr, "  -r role      User role.\n");
        fprintf(stderr, "  -t timeout   User session timeout.\n");
        fprintf(stderr, "  -v version   API version number. Default is 1.0\n");
        fprintf(stderr, "  -a method    one of md5, sha-1 or cleartext.\n");

        exit(1);
}

main(int argc, char **argv)
{
        char c, *p;
        int fd, len, postlen;
        struct sockaddr_in sa;

        while ((c = getopt(argc, argv, "a:i:m:n:p:r:t:v:")) != EOF) switch(c) {
                case 'i':       /* ipaddr */
                        ipaddr = optarg;
                        break;
                case 'm':       /* macaddr */
                        macaddr = optarg;
                        break;
                case 'n':       /* name */
                        name = optarg;
                        break;
                case 'p':       /* password */
                        password = optarg;
                        break;
                case 'r':       /* role */
                        role = optarg;
                        break;
                case 't':       /* session timeout */
                        tout = optarg;
                        break;
                case 'v':       /* version */
                        version = optarg;
                        break;
                case 'a':       /* authentication */
                        auth = optarg;
                        if (!strcasecmp(auth, "sha-1") &&
                                !strcasecmp(auth, "md5"))
                                usage();
                        break;
                default:
                        usage();
                        break;
        }
        argc -= (optind - 1);
        argv += (optind - 1);

        if ((argc < 3)) {
                usage();
        }
        if (version == NULL)
                version = "1.0";

        debug("server=%s, command=%s, version=%s, secret=%s\n",
                argv[1], argv[2], version, argv[3]?argv[3]:"<>");

        if (argv[3]) secret = argv[3];
```

```
        p = cmdbuf;
        sprintf(p, "xml=<aruba command='%s'>", argv[2]);
        p += strlen(p);
        if (ipaddr) {
                sprintf(p, "<ipaddr>%s</ipaddr>", ipaddr);
                p += strlen(p);
        }
        if (macaddr) {
                sprintf(p, "<macaddr>%s</macaddr>", macaddr);
                p += strlen(p);
        }
        if (name) {
                sprintf(p, "<name>%s</name>", name);
                p += strlen(p);
        }
        if (password) {
                sprintf(p, "<password>%s</password>", password);
                p += strlen(p);
        }
        if (role) {
                sprintf(p, "<role>%s</role>", role);
                p += strlen(p);
        }
        if (tout) {
                sprintf(p, "<session timeout>%s</session timeout>", tout);
                p += strlen(p);
        }
        if (secret) {
                if (auth == NULL) {
                        key = secret;
                        auth = "cleartext";
#ifndef OPENSSL_NO_SHA1
                } else if (!strcasecmp(auth, "sha-1")) {
                        key = enchashbuf;
                        SHA1(secret, strlen(secret), hashbuf);
                        encode_message_digest(hashbuf, 20, enchashbuf);
#endif
                } else if (!strcasecmp(auth, "md5")) {
                        key = enchashbuf;
                        md5_calc(hashbuf, secret, strlen(secret));
                        encode_message_digest(hashbuf, 16, enchashbuf);
                }
                debug("Message authentication is %s (%s)\n", auth, key);
                sprintf(p, "<authentication>%s</authentication><key>%s</key>",
                        auth, key);
                p += strlen(p);
        }
        debug("\n");
        sprintf(p, "<version>%s</version>", version);
        sprintf(p, "</authresponse>");
        cgi_escape_url(encbuf, sizeof(encbuf), cmdbuf, strlen(cmdbuf), 0);

        postlen = sprintf(post,
                "POST /auth/command.xml HTTP/1.0\r\n"
                "User-Agent: ecp\r\n"
                "Host: %s\r\n"
                "Pragma: no-cache\r\n"
                "Content-Length: %d\r\n"
                /* "Content-Type: application/x-www-form-urlencoded\r\n" */
                "Content-Type: application/xml\r\n"
                "\r\n"
```

```
                "%s",
                argv[1], strlen(encbuf), encbuf);

        inet_aton(argv[1], &sa.sin_addr);
        sa.sin_family = AF_INET;
        sa.sin_port = htons(80);
        fd = socket(AF_INET, SOCK_STREAM, 0);
        if (fd < 0) {
                perror("socket");
                exit(1);
        }
        if (connect(fd, (struct sockaddr *) &sa, sizeof(sa)) < 0) {
                perror("connect");
                exit(1);
        }

        if (write(fd, post, postlen) != postlen) {
                perror("write");
                exit(1);
        }

        while ((len = read(fd, post, sizeof(post))) > 0)
                write(1, post, len);
        close(fd);
        exit(0);
}

static void encode_message_digest (unsigned char *md, int mdlen, char *output)
{
        int i;

        for (i=0; i<mdlen; i++) {
                sprintf(output, "%02x", md[i]);
                output += 2;
        }
}

}
##### Authentication Script Example -- END --
```

Understanding Request and Response

The controller processes the authentication task and sends a response to the authentication server in the XML format to the authentication server. The XML response contains the status of the request and a code in case of an error. The example script is listed in Figure 243.

Request format: `<script_name> [options] <controller-ip> <command> <secret_key>`

Understanding XML API Request Parameters

The Table 227list all parameter that you can use in a request.

Table 227: *XML API Request Parameters and Descriptions*

Parameter	Description
`script_name`	The name of the script executable.
`Options`	• `-i <ip_addr>`—Specify the client's IP address. • `-m <mac_addr>`—Specify the client's MAC address. • `-n <name>`—Specify the client's user name. • `-p <passwd>`—Specify the client password. • `-r role`—Specify the current user role of the client. • `-t timeout`—User session timeout. • `-v version`—API version number. Default is 1.0 • `-a method`—Specify the encryption method to send the secret key. You can specify MD5 or SHA-1 or cleartext as the encryption method. By default, cleartext method is used to send the key. • `-s sessid`—Active session Id
`controller-ip`	The IP address of the controller that will receive the authentication requests.
command	The authentication command sent to the controller. You can send one of the following commands per request: • `add`: Adds the client to your network. • `delete`: Deletes the client from your network • `query`: Fetches information about the client • `blacklist`: Blacklists or block the client from connecting to your network • `authenticate`: Authenticates the client and assigns the default authenticated role.
`secret_key`	The password used to validate the authentication request from your authentication server. See Configuring the XML API Server on page 1024 for more information.

Understanding XMl API Response

The response message from the controller is sent in an XML format. The default format of the response is:

```
[Message header]
Displays the request parameters and other standard header details.
..
...
..

<response>
        <status>Status Message</status>
        <code>Code in case of an error</code>
</response>
```

Adding a Client

This command will add a client on your network.

Figure 244 *Adding a client—request and response*

```
john@linux:/home/john/tools/xml-api# ./auth -i 10.10.10.249 -m 00:19:d2:01:0b:aa -r logon
10.11.12.13 add $abcd$1234$
```

The commands sends the following information in the authentication request to the controller:

- Client IP address: **`10.10.10.249`**
- Client MAC address: **`00:19:d2:01:0b:aa`**
- Authentication server IP address: **`10.11.12.13`**
- Authentication command: **`add`**

- Key to validate authentication request: **$abcd$1234$**
- Verification key is sent in cleartext format

Response from the controller

```
server=10.11.12.13, command=add, version=1.0, secret=$abcd$1234$ sessid=
Message authentication is cleartext ($abcd$1234$)

HTTP/1.1 200 OK
Date: Tue, 03 Aug 2010 23:32:16 GMT
Server:
Connection: close
Content-Type: text/xml

<authresponse>
  <status>Ok</status>
  <code>0</code>
</authresponse>
```

View the updated details of the client on the controller

```
(host) #show user-table

Users
-----
    IP              MAC             Name      Role      Age(d:h:m)  Auth  ......
----------    ------------       ------    ----      ----------  ----  ... [truncated]
10.10.10.249  00:19:d2:01:0b:aa             logon     00:00:00          ......

User Entries: 1/1
```

Deleting a Client

This command will delete a client from your network.Deleting a client—request and response

```
john@linux:/home/john/tools/xml-api# ./auth -i 10.10.10.248 10.11.12.13 delete $abcd$1234$
```

This command sends the following information in the request to the controller:

- Client IP address: **10.10.10.248**
- Authentication server IP address: **10.11.12.13**
- Authentication command: **delete**
- Key to validate authentication request: **$abcd$1234$**
- Key is sent in cleartext format

Response from the controller

```
server=10.11.12.13, command=delete, version=1.0, secret=$abcd$1234$ sessid=
Message authentication is cleartext ($abcd$1234$)

HTTP/1.1 200 OK
Date: Tue, 03 Aug 2010 23:30:32 GMT
Server:
Content-Length: 56
Connection: close
Content-Type: text/xml

<authresponse>
  <status>Ok</status>
  <code>0</code>
</authresponse>
```

Authenticating a Client

This command will authenticate and change the role of a client. To illustrate the authentication command request process this section displays status of the client before and after the authentication command request.

Status of the client before authentication

The following `show user` command shows the role of the client is `logon` before the authentication request is processed by the controller.

```
(host) #show user

Users
-----
    IP              MAC             Name      Role      Age(d:h:m)  Auth  .....
----------    ------------       ------    ----      ----------  ----  ... [truncated]
10.10.10.248  00:19:d2:01:0b:84             logon     00:00:00          .....

User Entries: 1/1
```

The following command shows the captive portal status of the logon role of the client.

```
(host) (config-role) #show rights logon | include "Captive Portal profile"
 Captive Portal profile = default
```

Sending the authentication command

Use the `authenticate` keyword in the script to send the authentication command request.

Figure 245 *Authenticating the client—request and response*

```
john@linux:/home/john/tools/xml-api# ./auth -i 10.10.10.248 -n john -p password 10.11.23.24
authenticate $abcd$1234$
```

This commands sends the following information in the request to the controller:

- Client IP address: **`10.10.10.248`**
- Client username: **`john`**
- Client password: **`password`**
- Authentication server IP address: **`10.11.12.13`**
- Authentication command: **`authenticate`**
- Key to validate authentication request: **`$abcd$1234$`**
- Key is sent in cleartext format

Response from the controller

```
server=10.11.12.13, command=authenticate, version=1.0, secret=$abcd$1234$ sessid=
Message authentication is cleartext ($abcd$1234$)

HTTP/1.1 200 OK
Date: Tue, 03 Aug 2010 23:23:42 GMT
Server:
Connection: close
Content-Type: text/xml

<authresponse>
  <status>Ok</status>
  <code>0</code>
</authresponse>
```

Status of the client after authentication

The following `show user` command shows the role of the client is change to `guest` after the authentication request is processed by the controller.

```
(host) (config) #show user
Users
-----
    IP              MAC             Name   Role      Age(d:h:m)  Auth  .....
----------    ------------       ------  ----      ----------  ----  ... [truncated] .
10.10.10.248  00:19:d2:01:0b:84  John    guest     00:00:04    Web   .....

User Entries: 1/1
```

Querying for Client Details

This command will fetch a all details about a client connected in your network. Querying Client Information—request and response

```
john@linux:/home/john/tools/xml-api# ./auth -i 10.10.10.249 10.11.12.13 query $abcd$1234$
```

This commands sends the following information in the request to the controller:

- Client IP address: **10.10.10.249**
- Client username: **john**
- Client password: **password**
- Authentication server IP address: **10.11.12.13**
- Authentication command: **query**
- Key to validate authentication request: **$abcd$1234$**
- Key is sent in cleartext format

Response from the controller

```
server=10.11.12.13, command=query, version=1.0, secret=$abcd$1234$ sessid=
Message authentication is cleartext ($abcd$1234$)

HTTP/1.1 200 OK
Date: Tue, 03 Aug 2010 23:34:30 GMT
Server:
Connection: close
Content-Type: text/xml

<authresponse>
  <status>Ok</status>
  <code>0</code>
  <macaddr>00:19:d2:01:0b:aa</macaddr>
  <name>john</name>
  <role>logon</role>
  <type>Wireless</type>
  <vlan>1</vlan>
  <location>N/A</location>
  <age>00:00:02</age>
  <auth_status>Unauthenticated</auth_status>
  <essid></essid>
  <bssid>00:00:00:00:00:00</bssid>
  <phy_type>b</phy_type>
  <mobility_state>Wireless</mobility_state>
  <in_packets>0</in_packets>
  <in_octets>0</in_octets>
  <out_packets>0</out_packets>
  <out_octets>0</out_octets>
```

```
</authresponse>
```

The output of the `show user` command displays the client information.

```
Users
-----
    IP              MAC             Name     Role      Age(d:h:m)  Auth  .....
----------    ------------       ------   ----      ----------  ----  ... [truncated]
10.10.10.249  00:19:d2:01:0b:aa  John     logon     00:00:01          .....

User Entries: 1/1
```

Blacklisting a Client

This command will blacklist a client and restrict it from connecting to your network. The `show user-table` lists the client connected on your network before processing the request to blacklist the client.

```
Users
-----
    IP              MAC             Name     Role      Age(d:h:m)  .....
----------    ------------       ------   ----      ----------  ... [truncated] ...
10.10.10.248  00:19:d2:01:0b:84  John     guest     00:00:00    .....

User Entries: 1/1
```

Figure 246 *Blacklisting a Client—request and response*

```
john@linux:/home/john/tools/xml-api# ./auth -i 10.10.10.248 10.11.12.13 blacklist $abcd$1234$
```

This commands sends the following information in the request to the controller:

- Client IP address: **`10.10.10.248`**
- Authentication server IP address: **`10.11.12.13`**
- Authentication command: **`blacklist`**
- Key to validate authentication request: **`$abcd$1234$`**
- Key is sent in cleartext format

Response from the controller

```
server=10.11.12.13, command=blacklist, version=1.0, secret=$abcd$1234$ sessid=
Message authentication is cleartext ($abcd$1234$)

HTTP/1.1 200 OK
Date: Tue, 03 Aug 2010 23:29:11 GMT
Server:
Content-Length: 56
Connection: close
Content-Type: text/xml

<authresponse>
  <status>Ok</status>
  <code>0</code>
</authresponse>
```

The `show user-table` command does not list the blacklisted client. You can use the show ap blacklist-clients command on your controller to view the list of blacklisted clients

```
(host) (config) #show ap blacklist-clients

Blacklisted Clients
-------------------
STA                reason        block-time(sec)  remaining time(sec)
---                ------        ---------------  -------------------
00:19:d2:01:0b:84  user-defined  5                3595
```

Topics in this chapter include:

- Understanding Mode Support on page 1039
- Understanding Basic System Defaults on page 1040
- Understanding Default Management User Roles on page 1049
- Understanding Default Open Ports on page 1052

Understanding Mode Support

Most ArubaOS features are supported in all forwarding modes. However, there are a some features that are not supported in one or more forwarding modes. Campus APs do not support split-tunnel forwarding mode and the decrypt-tunnel forwarding mode does not support TKIP Counter measure management on campus APs or remote APs.

Table 228 describes the features that are not supported in each forwarding mode.

Table 228: *Features not Supported in Each Forwarding Mode*

Forwarding Mode	Feature Not Supported
Split Tunnel Mode on Remote APs	VLAN Pooling Named VLAN Voice over Mesh Video over Mesh Layer-2 Mobility Layer-3 Mobility IGMP Proxy Mobility Mobile IP TKIP countermeasure mgmt Bandwidth based CAC Dynamic Multicast Optimization
Bridge Mode on Campus APs or Remote APs	Firewall—SIP/SCCP/RTP/RTSP Voice Support Firewall—Alcatel NOE Support Voice over Mesh Video over Mesh Named VLAN Captive portal Rate Limiting for broadcast/multicast Power save: Wireless battery boost Power save: Drop wireless multicast traffic Power save: Proxy ARP (global) Power save: Proxy ARP (per-SSID) Automatic Voice Flow Classification
Bridge Mode on Campus APs or Remote APs (continued)	SIP ALG SIP: SIP authentication tracking SIP: CAC enforcement enhancements SIP: Phone number awareness SIP: R-Value computation SIP: Delay measurement

Forwarding Mode	Feature Not Supported
	Management: Voice-specific views Management: Voice client statistics Management: Voice client troubleshooting Voice protocol monitoring/reporting SVP ALG H.323 ALG Vocera ALG SCCP ALG NOE ALG Layer 3 Mobility IGMP Proxy Mobility Mobile IP TKIP countermeasure mgmt Bandwidth based CAC Dynamic Multicast Optimization

Understanding Basic System Defaults

The default administrator user name is **admin**, and the default password is also **admin**.The ArubaOS software includes several predefined network services, firewall policies, and roles.

Network Services

Table 229 lists the predefined network services and their protocols and ports.

Table 229: *Predefined Network Services*

Name	Protocol	Port(s)
svc-dhcp	udp	67 68
svc-snmp-trap	udp	162
svc-smb-tcp	tcp	445
svc-https	tcp	443
svc-ike	udp	500
svc-l2tp	udp	1701
svc-syslog	udp	514
svc-pptp	tcp	1723
svc-telnet	tcp	23
svc-sccp	tcp	2000
svc-tftp	udp	69
svc-sip-tcp	tcp	5060
svc-kerberos	udp	88

Name	Protocol	Port(s)
svc-pop3	tcp	110
svc-adp	udp	8200
svc-noe	udp	32512
svc-noe-oxo	udp	5000
svc-dns	udp	53
svc-msrpc-tcp	tcp	135 139
svc-rtsp	tcp	554
svc-http	tcp	80
svc-vocera	udp	5002
svc-nterm	tcp	1026 1028
svc-sip-udp	udp	5060
svc-papi	udp	8211
svc-ftp	tcp	21
svc-natt	udp	4500
svc-svp	119	0
svc-gre	gre	0
svc-smtp	tcp	25
svc-smb-udp	udp	445
svc-esp	esp	0
svc-bootp	udp	67 69
svc-snmp	udp	161
svc-icmp	icmp	0
svc-ntp	udp	123
svc-msrpc-udp	udp	135 139
svc-ssh	tcp	22
svc-h323-tcp	tcp	1720
svc-h323-udp	udp	1718 1719
svc-http-proxy1	tcp	3128

Name	Protocol	Port(s)
svc-http-proxy2	tcp	8080
svc-http-proxy3	tcp	8888
svc-sips	tcp	5061
svc-v6-dhcp	udp	546 547
svc-v6-icmp	icmp	0
any	any	0

Policies

The following are predefined policies.

Table 230: *Predefined Policies*

Predefined Policy	Description
ip access-list session allowall `any any any permit`	An "allow all" firewall rule that permits all traffic.

`ip access-list session control`
`user any udp 68 deny`
`any any svc-icmp permit`
`any any svc-dns permit`
`any any svc-papi permit`
`any any svc-cfgm-tcp permit`
`any any svc-adp permit`
`any any svc-tftp permit`
`any any svc-dhcp permit`
`any any svc-natt permit`
	Controls traffic—Apply to untrusted wired ports in order to allow Dell APs to boot up.
	NOTE: In most cases wired ports should be made "trusted" when attached to an internal network.

`ip access-list session captiveportal`
`user alias mswitch svc-https dst-nat 8081`
`user any svc-http dst-nat 8080`
`user any svc-https dst-nat 8081`
`user any svc-http-proxy1 dst-nat 8088`
`user any svc-http-proxy2 dst-nat 8088`
`user any svc-http-proxy3 dst-nat 8088`
	Enables Captive Portal authentication.
	1. Any HTTPS traffic destined for the controller will be NATed to port 8081, where the captive portal server will answer.
	2. All HTTP traffic to any destination will be NATed to the controller on port 8080, where an HTTP redirect will be issued.

Predefined Policy	Description
	3. All HTTPS traffic to any destination will be NATed to the controller on port 8081, where an HTTP redirect will be issued. 4. All HTTP proxy traffic will be NATed to the controller on port 8088. **NOTE:** In order for captive portal to work properly, DNS must also be permitted. This is normally done in the "logon-control" firewall rule.
`ip access-list session cplogout user alias mswitch svc-https dst-nat 8081`	Used to enable the captive portal "logout" window. If the user attempts to connect to the controller on the standard HTTPS port (443) the client will be NATed to port 8081, where the captive portal server will answer. If this rule is not present, a wireless client may be able to access the controller's administrative interface.

`ip access-list session vpnlogon`
` any any svc-ike permit`
` any any svc-esp permit`
` any any svc-l2tp permit`
` any any svc-pptp permit`
` any any svc-gre permit`
	This policy permits VPN sessions to be established to any destination. IPsec (IKE, ESP, and L2TP) and PPTP (PPTP and GRE) are supported.

`ip access-list session ap-acl`
`  any any udp 5000`
`  any any udp 5555`
`  any any svc-gre permit`
`  any any svc-syslog permit`
`  any user svc-snmp permit`
`  user any svc-snmp-trap permit`
`  user any svc-ntp permit`
	This is a policy for internal use and should not be modified. It permits APs to boot up and communicate with the controller.

Predefined Policy	Description
`ip access-list session validuser` `any any any permit`	This firewall rule controls which users will be added to the user-table of the controller through untrusted interfaces. Only IP addresses permitted by this ACL will be admitted to the system for further processing. If a client device attempts to use an IP address that is denied by this rule, the client device will be ignored by the controller and given no network access. You can use this rule to restrict foreign IP addresses from being added to the user-table. This policy should not be applied to any user role, it is an internal system policy.
ip access-list session vocera-acl any any svc-vocera permit queue high	Use for Vocera VoIP devices to automatically permit and prioritize Vocera traffic.
ip access-list session icmp-acl `any any svc-icmp permit`	Permits all ICMP traffic.
ip access-list session sip-acl any any svc-sip-udp permit queue high any any svc-sip-tcp permit queue high	Use for SIP VoIP devices to automatically permit and prioritize all SIP control and data traffic.
ip access-list session https-acl any any svc-https permit	Permits all HTTPS traffic.
ip access-list session dns-acl any any svc-dns permit	Permits all DNS traffic.

Predefined Policy	Description
ip access-list session logon-control user any udp 68 deny any any svc-icmp permit any any svc-dns permit any any svc-dhcp permit any any svc-natt permit	The default pre-authentication role that should be used by all wireless clients. Prohibits the client from acting as a DHCP server. Permits all ICMP, DNS, and DHCP. Also permits IPsec NAT-T (UDP 4500). Remove NAT-T if not needed.
ip access-list session srcnat user any any src-nat	This policy can be used to source-NAT all traffic. Because no NAT pool is specified, traffic that matches this policy will be source NATed to the IP address of the controller.
ip access-list session skinny-acl any any svc-sccp permit queue high	Use for Cisco Skinny VoIP devices to automatically permit and prioritize VoIP traffic.
ip access-list session tftp-acl any any svc-tftp permit	Permits all TFTP traffic.
ip access-list session guest	This policy is not used.
ip access-list session dhcp-acl any any svc-dhcp permit	Permits all DHCP traffic. If DHCP is not allowed, clients will not be able to request or renew IP addresses.
ip access-list session http-acl any any svc-http permit	Permits all HTTP traffic.
ip access-list session svp-acl any any svc-svp permit queue high user host 224.0.1.116 any permit	Use for Spectralink VoIP devices to automatically permit and prioritize Spectralink Voice Protocol (SVP).

Predefined Policy	Description
ip access-list session noe-acl any any svc-noe permit queue high	Use for Alcatel NOE VoIP devices to automatically permit and prioritize NOE traffic.
ip access-list session h323-acl any any svc-h323-tcp permit queue high any any svc-h323-udp permit queue high	Use for H.323 VoIP devices to automatically permit and prioritize H.323 traffic.
ipv6 access-list session v6-control user any udp 68 deny any any svc-v6-icmp permit any any svc-v6-dhcp permit any any svc-dns permit any any svc-tftp permit	Provides equivalent functionality to the "control" policy, but for IPv6 clients.
ipv6 access-list session v6-icmp-acl any any svc-v6-icmp permit	Permits all ICMPv6 traffic.
ipv6 access-list session v6-https-acl any any svc-https permit	Permits all IPv6 HTTPS traffic.
ipv6 access-list session v6-dhcp-acl any any svc-v6-dhcp permit	Permits all IPv6 DHCP traffic.
ipv6 access-list session v6-dns-acl any any svc-dns permit	Permits all IPv6 DNS traffic.
ipv6 access-list session v6-allowall any any any permit	Permits all IPv6 traffic.
ipv6 access-list session v6-http-acl any any svc-http permit	Permits all IPv6 HTTP traffic.
ipv6 access-list session v6-tftp-acl any any svc-tftp permit	Permits all IPv6 TFTP traffic.

ipv6 access-list session v6-logon-control
`user any udp 68 deny`
`any any svc-v6-icmp permit`
`any any svc-v6-dhcp permit`
`any any svc-dns permit`
	Provides equivalent functionality to the "logon-control" policy, but for IPv6 clients.

Validuser and Logon-control ACLs

Default firewall rules for both the validuser and logon-control ACLs prevent malicious users from ip spoofing source addresses the default firewall rule in the validuser ACL causes the packet to be dropped.

A client with the correct source address can send traffic to the below networks as a destination IP address. To deny traffic, the default firewall rule added to logon-control ACL denies traffic to the reserved addresses from user with the logon role.

The following networks can be blocked by the default firewall rules in both the validuser and logon-control ACLs:

- Network packets where the source address of the network packet is defined as being on a broadcast network (source address == 255.255.255.255)
- Network packets where the source address of the network packet is defined as being on a multicast network (source address = 224.0.0.0 – 239.255.255.255)
- Network packets where the source address of the network packet is defined as being a loopback address (127.0.0.1 through 127.255.255.254)
- Network packets where the source or destination address of the network packet is a link-local address (169.254.0.0/16)
- Network packets where the source or destination address of the network packet is defined as being an address “reserved for future use” as specified in RFC 5735 for IPv4; (240.0.0.0/4)
- Network packets where the source or destination address of the network packet is defined as an “unspecified address”(::/128) or an address “reserved for future definition and use”(addresses other than 2000::/3) as specified in RFC 3513 for IPv6. The IPv6 “an unspecified address”(::/128) is currently being checked in datapath and the packet is dropped. This is the default behavior and you can view the logs by enabling **firewall enable-per-packet-logging** configuration.

Roles

The following are predefined roles.



If you upgrade from a previous ArubaOS release, your existing configuration may have additional or different predefined roles. The information in this section only describes the predefined roles for this release.

Table 231: *Predefined Roles*

Predefined Role	Description
user-role ap-role session-acl control session-acl ap-acl	This is an internal role and should not be edited.
user-role default-vpn-role session-acl allowall ipv6 session-acl v6-allowall	This is the default role used for VPN-connected clients. It is referenced in the default "aaa authentication vpn" profile.
user-role voice session-acl sip-acl session-acl noe-acl session-acl svp-acl session-acl vocera-acl session-acl skinny-acl session-acl h323-acl session-acl dhcp-acl session-acl tftp-acl session-acl dns-acl session-acl icmp-acl	This role can be applied to voice devices in order to automatically permit and prioritize all VoIP protocols.
user-role guest session-acl http-acl session-acl https-acl session-acl dhcp-acl session-acl icmp-acl session-acl dns-acl ipv6 session-acl v6-http-acl	This is a default role for guest users. It permits only HTTP, HTTPS, DHCP, ICMP, and DNS for the guest user. To increase security, a "deny" rule for internal network destinations could be added at the beginning.

Predefined Role	Description
ipv6 session-acl v6-https-acl ipv6 session-acl v6-dhcp-acl ipv6 session-acl v6-icmp-acl ipv6 session-acl v6-dns-acl	
user-role guest-logon captive-portal default session-acl logon-control session-acl captiveportal	This role is used as the pre-authentication role for guest SSIDs. It allows control traffic such as DNS, DHCP, and ICMP, and also enables captive portal.
user-role <ssid>-guest-logon captive-portal default session-acl logon-control session-acl captiveportal	This role is only generated when creating a new WLAN using the WLAN Wizard. The WLAN Wizard creates this role when captive portal is enabled. This is the initial role that a guest will be placed in prior to captive portal authentication. By using a different guest logon role for each SSID, it is possible to enable multiple captive portal profiles with different customization.
user-role stateful-dot1x	This is an internal role used for Stateful 802.1x. It should not be edited.
user-role authenticated session-acl allowall ipv6 session-acl v6-allowall	This is a default role that can be used for authenticated users. It permits all IPv4 and IPv6 traffic for users who are part of this role.
user-role logon session-acl logon-control session-acl captiveportal session-acl vpnlogon ipv6 session-acl v6-logon-control	This is a system role that is normally applied to a user prior to authentication. This applies to wired users and non-802.1x wireless users. The role allows certain control protocols such as DNS, DHCP, and ICMP, and also enables captive portal and VPN termination/pass through. The logon role should be edited to provide only the required services to a pre-authenticated user. For example, VPN pass through should be disabled if it is not needed.
user-role <ssid>-logon session-acl control session-acl captiveportal session-acl vpnlogon	This role is only generated when creating a new WLAN using the WLAN Wizard. The WLAN Wizard creates this role when captive portal is enabled and a PEFNG license is installed. This is the initial role that a client will be placed in prior to captive portal authentication. By using a different logon role for each SSID, it is possible to enable multiple captive portal profiles with different customization.
user-role <ssid>-captiveportal-profile	When utilizing the WLAN Wizard and you do not have a PEF NG installed and you are configuring an Internal or Guest WLAN with captive portal enabled, the controller creates an implicit user role with the same name as the captive portal profile, <ssid>-captiveportal-profile. This implicit user role allows only DNS and DHCP traffic between the client and network and directs all HTTP or HTTPS requests to the captive portal. You cannot directly modify the implicit user role or its rules. Upon authentication, captive portal clients are allowed full access to their assigned VLAN. Once the WLAN configuration is pushed to the controller, the WLAN wizard will associate the new role with the initial user role that you specify in the AAA profile. This role will not be visible to the user in the WLAN wizard.

Understanding Default Management User Roles

The ArubaOS software includes predefined management user roles.



If you upgrade from a previous ArubaOS release, your existing configuration may have different management roles. The information in this section only describes the predefined management roles for this release.

Table 232: *Predefined Management Roles*

Predefined Role	Permissions
root	This role permits access to all management functions (commands and operations) on the controller.
read-only	This role permits access to CLI `show` commands or WebUI monitoring pages only.
guest-provisioning	This role permits access to configuring guest users in the controller's internal database only. This user only has access via the WebUI to create guest accounts; there is no CLI access. Guest-provisioning tasks include creating or generating the user name and password for a guest account as well as configuring when the account expires.
location-api-mgmt	This role permits access to location API information and the CLI; however, you cannot use any CLI commands. This role does not permit access to the WebUI. Using a third-party location appliance, you can gather information about the location of 802.11 stations. To log in to the controller using a third-party location appliance, enter: http[s]://<ipaddress>[:port]/screens/wms/wms.login. You are prompted to enter your username and password (for example, the username and password associated with the location API management role). Once authenticated, you can use an API call to request location information from the controller, for example: http[s]://<ipaddress>[:port]/screens/wms/wms.cgi?opcode=wlm-get-spot&campus-name=<campus id>&building-name<building id>&mac=<client1>,<client2>....
network-operations	This role supports a subset of show, configuration, action, and database commands that are used to monitor the controller. You can log into the CLI; however, you can only use a subset of CLI commands to monitor the controller. This role permits the following WebUI pages and associated CLI commands: As a network-operations user, commands with an asterisk (*) are hidden in the CLI but are executed and visible from the WebUI. Plan Page • You can move APs on the floor plan and save their new location. • You cannot change or modify the AP configuration. Reports Page • You can view all of the available reports. Events Page • You can view all of the available events. Monitoring Page You can view the reports created by the following CLI commands: • show keys all • show mobility-managers • show roleinfo • show license* • show ap essid • DB:opcode=cr-load Monitoring > Network > Network Summary You can view the reports created by the following CLI commands:

Predefined Role	Permissions
	• show interface vlan <id> • show interface loopback • show datapath utilization • show aaa state configuration • show user-table unique • show aaa authentication-server all • show switches summary • show ap blacklist-clients • show wlan-ap-count type access-points* • show wlan-ap-count type air-monitor* • show wlan-ap-count type secure-access* • show user-table verbose • show ap database unprovisioned page <page> • show ap-group default • show wlan virtual-ap • show rf dot11a-radio-profile • show rf dot11g-radio-profile • show ap wired-ap-profile • show ap enet-link-profile • show ap system-profile • show wlan voip-cac-profile • show wlan traffic-management-profile • show ap regulatory-domain-profile • show ap snmp-profile • show rf optimization-profile • show rf event-thresholds-profile • show ids profile • show rf arm-profile • show ap association bssid

Predefined Role	Permissions
network-operations (continued)	**Monitoring** > **Network** > **All Access PointsMonitoring** > **Network** > **All Wired Access Points** You can view the reports created by the following CLI commands: • `DB:opcode=monitor-summary` • `DB:opcode=cr-load` • `DB:opcode=wlm-search&class=probes&start` • `DB:opcode=wlm-search&class=amii` • `DB:opcode=monitor-get-all-gps&status=any` • `show ap-group` • `show vlan status` **Monitoring** > **Controller** > **Controller Summary** You can view the reports created by the following CLI commands: • `show switches` • `show switches summary` **Monitoring** > **Controller** > **Air Monitors** You can view the reports created by the following CLI commands: • `show wlan-ap start*` **Monitoring** > **Controller** > **Clients** You can view the reports created by the following CLI commands: • `show ip mobile host` • `show ip mobile trail {<ipaddr> | <macaddr>}` • `show esi groups` • `show esi servers` • `show esi ping` • `show esi parser stats` • `show private port status*` • `show vlan` • `show port stats` • `show spanning-tree interface fastethernet <slot/port>` • `show interface fastethernet <slot/port> counters` • `clear counters fastethernet <slot/port>` • `show snmp trap-queue <page>` **Monitoring** > **Controller** > **Clients** > **Packet CaptureMonitoring** >**Controller** > **Clients** > **LocateMonitoring** > **Controller** > **Clients** > **Debug** You can view the reports created by the following CLI commands: • `aaa user debug mac` **Monitoring** > **Controller**> **Clients** > **Disconnect** You can view the reports created by the following CLI commands: • `stm kick-off-sta <macaddr>` • `aaa user logout <ipaddr>`
network-operations (continued)	**Monitoring** > **Controller**> **Clients** > **Blacklist** You can view the reports created by the following CLI commands: • `stm add-blacklist-client <macaddr>` • `aaa user delete {<ipaddr> | all | mac <macaddr> | name <username> | role <role>}` **Monitoring** > **Controller** > **Blacklist Clients** You can view the reports created by the following CLI commands: • `stm remove-blacklist-client <macaddr>` **Monitoring** > **Controller** > **External Services Interface** You can view the reports created by the following CLI commands: • `show esi groups`

Predefined Role	Permissions
	• `show esi servers` • `show esi ping` • `show esi parser stats` **Monitoring** > **Controller** > **Ports** You can view the reports created by the following CLI commands: • `show model-switch-internal*` • `show slots` • `show private port status*` • `show vlan` **Monitoring** > **Controller**> **Inventory** You can view the reports created by the following CLI commands: • `show keys` **Monitoring** > **WLAN** You can view the reports created by the following CLI commands: • `DB:opcode=get-permissions` • `DB:opcode=cr-load` • `show switches` • `show switches summary` **Monitoring** > **Voice** You can view the reports created by the following CLI commands: • `show ap association voip-only` • `show ap active voip-only` • `show voice call-counters` • `show voice client status` • `show voice call-quality` • `show voice call-density` • `show voice call-cdrs` • `show voice call-perf`

Understanding Default Open Ports

By default, Dell controllers and access points treat ports as untrusted. However, certain ports are open by default only on the trusted side of the network. These open ports are listed in Table 233.

Table 233: *Default (Trusted) Open Ports*

Port Number	Protocol	Where Used	Description
17	TCP	controller	This is use for certain types of VPN clients that accept a banner (QOTD). During normal operation, this port will only accept a connection and immediately close it.
21	TCP	controller	
22	TCP	controller	SSH
23	TCP	AP and controller	Telnet is disabled by default but the port is still open.
53	UDP	controller	Internal domain.

Port Number	Protocol	Where Used	Description
67	UDP	AP (and controller if DHCP server is configured)	DHCP server.
68	UDP	AP (and controller if DHCP server is configured)	DHCP client.
69	UDP	controller	TFTP
80	TCP	AP and controller	HTTP Used for remote packet capture where the capture is saved on the Access Point. Provides access to the WebUI on the controller.
123	UDP	controller	NTP
161	UDP	AP and controller	SNMP. Disabled by default.
443	TCP	controller	Used internally for captive portal authentication (HTTPS) and is exposed to wireless users. A default self-signed certificate is installed in the controller. Users in a production environment are urged to install a certificate from a well known CA such as Verisign. Self-signed certs are open to man-in-the-middle attacks and should only be used for testing. **Required for VIA**: During the initializing phase, VIA uses HTTPS connections to perform trusted network and captive portal checks against the controller. It is mandatory that you enable port 443 on your network to allow VIA to perform these checks
500	UDP	controller	ISAKMP
514	UDP	controller	Syslog
1701	UDP	controller	L2TP
1723	TCP	controller	PPTP
2300	TCP	controller	Internal terminal server opened by `telnet soe` command.
3306	TCP	controller	Remote wired MAC lookup.
4343, 443	TCP	controller	HTTPS.Both port 4343 and 443 are supported. If port 4343 is used it redirects to port 443. If port 443 is used it continues to connect using this port. A default self-signed certificate is installed in the controller. Users in a production environment are urged to install a certificate from a well known CA such as Verisign. Self-signed certs are open to man-in-the-middle attacks and should only be used for testing
4500	UDP	controller	sae-urn

Port Number	Protocol	Where Used	Description
			Required for VIA: During the initializing phase, VIA uses HTTPS connections to perform trusted network and captive portal checks against the controller. It is mandatory that you enable port 4500 on your network to allow VIA to perform these checks
8080	TCP	controller	Used internally for captive portal authentication (HTTP-proxy). This port is not exposed to wireless users.
8081	TCP	controller	Used internally for captive portal authentication (HTTPS). Not exposed to wireless users. A default self-signed certificate is installed in the controller. Users in a production environment are urged to install a certificate from a well known CA such as Verisign. Self-signed certs are open to man-in-the-middle attacks and should only be used for testing.
8082	TCP	controller	Used internally for single sign-on authentication (HTTP). Not exposed to wireless users.
8083	TCP	controller	Used internally for single sign-on authentication (HTTPS). Not exposed to wireless users.
8088	TCP	controller	For internal use.
8200	UDP	controller	The Aruba Discovery Protocol (ADP)
8211	UDP	controller	For internal use.
8888	TCP	controller	Used for HTTP access.

This chapter describes how to configure several DHCP vendor-specific options.

Topics in this chapter include:

- Configuring a Windows-Based DHCP Server on page 1055
- Enabling DHCP Relay Agent Information Option (Option 82) on page 1058
- Enabling Linux DHCP Servers on page 1059

Configuring a Windows-Based DHCP Server

Configuring a Microsoft Windows-based DHCP server to send option 43 to the DHCP client on a Dell AP consists of the following two tasks:

- Configuring Option 60
- Configuring Option 43

DHCP servers are a popular way of configuring clients with basic networking information such as an IP address, a default gateway, network mask, DNS server, and so on. Most DHCP servers have the ability to also send a variety of optional information, including the Vendor-Specific Option Code, also called option 43.

When a client or an AP requests for option 43 (Vendor Specific Information), the controller responds with the value configured by administrator in the DHCP pool.

Configuring Option 60

This section describes how to configure the Vendor Class Identifier Code (option 60) on a Microsoft Windows-based DHCP server.

As mentioned in the overview section, option 60 identifies and associates a DHCP client with a particular vendor. Any DHCP server configured to take action based on a client's vendor ID should also have this option configured.

Since option 60 is not a predefined option on a Windows DHCP server, you must add it to the option list for the server.

To configure option 60 on the Windows DHCP server

1. On the DHCP server, open the DHCP server administration tool by clicking **Start** > **AdministrativeTools** > **DHCP**.
2. Find your server and right-click on the scope to be configured under the server name. Select **Set Predefined Options**.
3. In the Predefined Options and Values dialog box, click **Add.**
4. In the Option Type dialog box, enter the following information

Table 234: *Configure option 60 on the Windows DHCP server*

Field	Information
Name	Dell Access Point
Data Type	String
Code	60
Description	Dell AP vendor class identifier

5. Click **OK** to save this information.
6. In the Predefined Options and Values dialog box, make sure **060 Dell Access Point** is selected from the Option Name drop-down list.
7. In the Value field, enter the following information:

   String : DellAP
8. Click **OK** to save this information.
9. Under the server, select the scope you want to configure and expand it. Select **Scope Options**, then select **Configure Options**.
10. In the Scope Options dialog box, scroll down and select **060 Dell Access Point**. Confirm the value is set to **DellAP** and click **OK**.
11. Confirm that the option **060 Dell Access Point** is listed in the right pane.

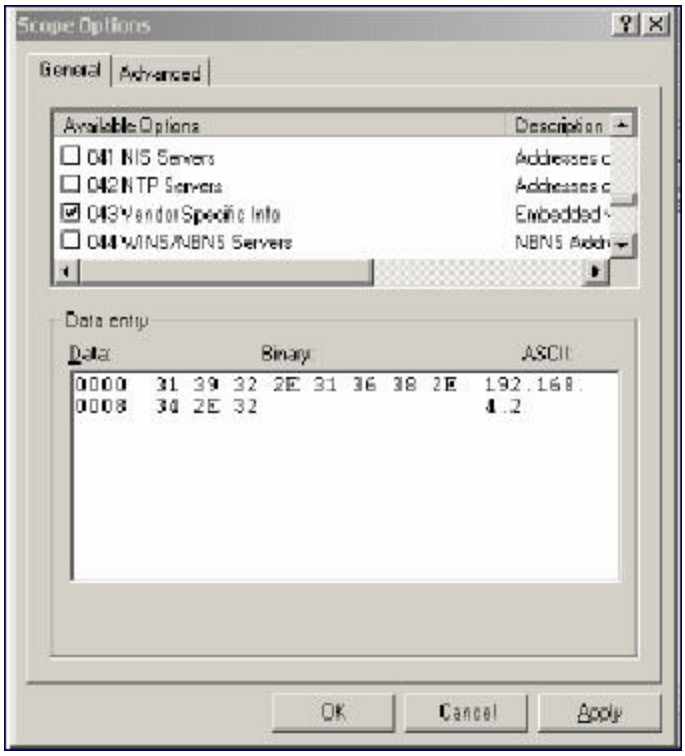
Configuring Option 43

Configuring Option 43 returns the IP address of the Dell master controller to a Dell DHCP client. This information allows Dell APs to auto-discover the master controller and obtain their configuration.

To configure option 43 on the Windows DHCP server:

1. On the DHCP server, open the DHCP server administration tool by clicking **Start > Administration Tools > DHCP**.
2. Find your server and right-click on the scope to be configured under the server name. Click on the Scope Options entry and select **Configure Options**.
3. In the Scope Options dialog box (Figure 247), scroll down and select **043 Vendor Specific Info**.

Figure 247 *Scope Options Dialog Box.*



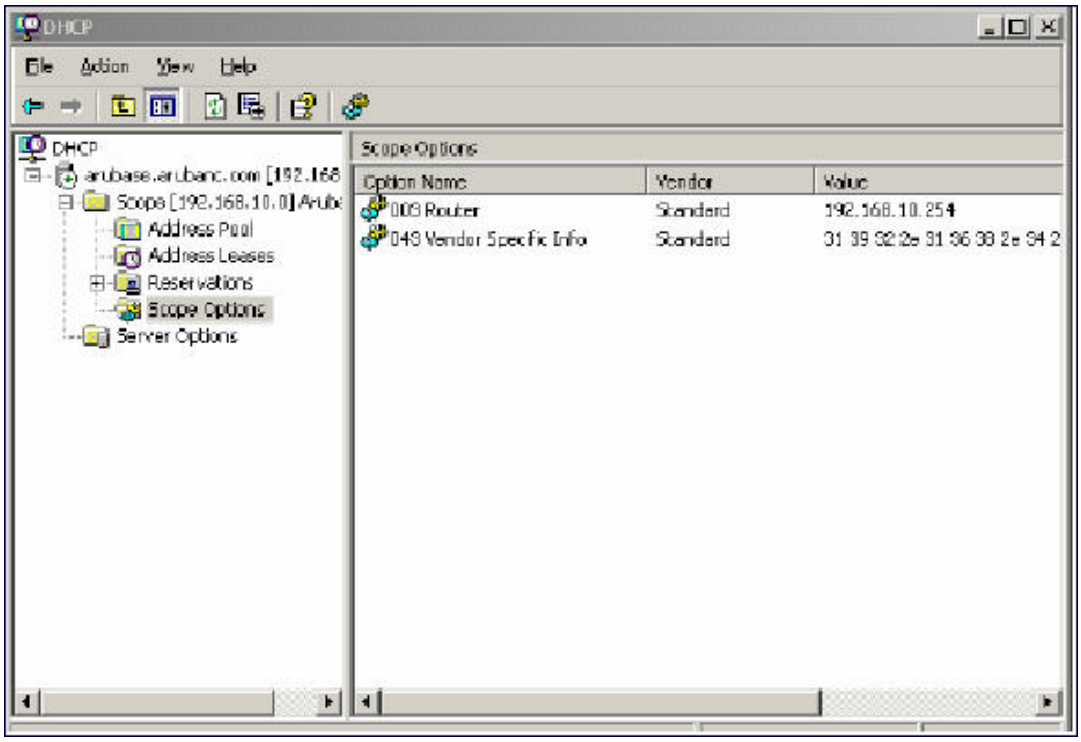
4. In the Data Entry field, click anywhere in the area under the ASCII heading and enter the following information:

   ASCII : Loopback address of the master controller

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

Option 43 is configured for this DHCP scope. Note that even though you entered the IP address in ASCII text, it displays in binary form.

Figure 248 *DHCP Scope Values*



Enabling DHCP Relay Agent Information Option (Option 82)

The DHCP Relay Agent Information option (Option 82) allows the DHCP Relay Agent to insert circuit specific information into a request that is being forwarded to a DHCP server.

The controller, when acting as a DHCP relay agent, inserts information about the AP and SSID through which a client is connecting into the DHCP request. Many service providers use this mechanism to make access control decisions.

Configuring Option 82

You can configure Option 82 using the WebUI or the CLI. You can include only the MAC address or MAC address and ESSID. The MAC address is the hardware address and ESSID is an alphanumeric name that uniquely identifies a wireless network.

In the WebUI

1. Navigate to **Configuration** > **Network** > **IP** > **IP Interfaces**.
2. Click **Edit** next to the VLAN ID for which you want to configure Option 82.
3. Under DHCP Helper Address select **Mac** or **Mac Essid** from the Option-82 drop-down menu.
4. Click **Apply**.

In the CLI

This example enables Option 82 for VLAN 5 using ESSID. You can include only the MAC address or MAC address and ESSID.

```
(host)(config)#interface vlan 5
(host)(config-subif)#option-82
(host)(config-subif)#option-82 mac essid
```

```
(host)(config-subif)#
```

Enabling Linux DHCP Servers

The following is an example configuration for the Linux dhcpd.conf file. After you enter the configuration, you must restart the DHCP service.

```
option serverip code 43 = ip-address;
class "vendor-class" {
      match option vendor-class-identifier;
}
.
.
.
subnet 10.200.10.0 netmask 255.255.255.0 {
   default-lease-time 200;
   max-lease-time 200;
   option subnet-mask 255.255.255.0;
   option routers 10.200.10.1;
   option domain-name-servers 10.4.0.12;
   option domain-name "vlan10.aa.mycorpnetworks.com";
   subclass "vendor-class" "ArubaAP" {
      option vendor-class-identifier "ArubaAP";
#
# option serverip <loopback-IP-address-of-master-controller>
#
      option serverip 10.200.10.10;
   }
   range 10.200.10.200 10.200.10.252;
}
```

This chapter provides examples of how to configure a Microsoft Internet Authentication Server, and a Windows XP wireless client for 802.1X authentication with the controller (see 802.1X Authentication on page 253). for information about configuring the controller

For more information about configuring computers in a Windows environment for PEAP-MS-CHAPv2 and EAP-TLS authentication, see the Microsoft document *Step-by-Step Guide for Setting Up Secure Wireless Access in a Test Lab*, available from Microsoft's Download Center (at microsoft.com/downloads. Additional information on client configuration is available at microsoft.com/technet/prodtechnol/winxppro/maintain/wificomp.mspx#EQGAC.

This chapter describes the following topics:

- Configuring Microsoft IAS on page 1060
- Configuring Management Authentication using IAS on page 1062
- Window XP Wireless Client Sample Configuration on page 1064

Configuring Microsoft IAS

Microsoft Internet Authentication Server (IAS) provides authentication functions for the wireless network. IAS implements the RADIUS protocol, which is used between the Dell controller and the server. IAS uses Active Directory as the database for looking up computers, users, passwords, and group information.

RADIUS Client Configuration

Each device in the network that needs to authenticate to a RADIUS server must be configured as a RADIUS client. You must configure the Dell controller as a RADIUS client.



The steps to perform this task may very depending on the version of Windows currently running on your server. For complete details on configuring Windows IAS, refer to the Windows documentation available (at microsoft.com/downloads).

To configure a RADIUS client:

1. From your windows server, navigate to **Start > Settings > Control Panel > Administrative Tools>Internet Authentication Service**.
2. In the Internet Authentication Service window, select **RADIUS Clients**.
3. To configure a RADIUS client, select **Action > New RADIUS Client** from the menu at the top of the window.
4. In the **New RADIUS Client** dialog window, enter the name and IP address for the controller. Click **Next**.
5. In the next window that appears, enter and confirm a shared secret. The shared secret is configured on both the RADIUS server and client, and ensures that an unauthorized client cannot perform authentication against the server.
6. Click **Finish**.

Remote Access Policies

The IAS policy configuration defines all policies related to wireless access, including time of day restrictions, session length, authentication type, and group-related policies. See Microsoft product documentation for

detailed descriptions and explanations of IAS policy settings.

Active Directory Database

The Active Directory database serves as the master authentication database for both the wired and wireless networks. The IAS authentication server bases all authentication decisions on information in the Active Directory database. IAS is normally used as an authentication server for remote access and thus looks to the Active Directory "Remote Access" property to determine whether authentication requests should be allowed or denied. This property is set on a per-user or per-computer basis. For a user or computer to be allowed access to the wireless network, the remote access property must be set to "Allow access".

The authentication policy configured in IAS depends on the group membership of the computer or user in Active Directory. These policies are responsible for passing group information back to the controller for use in assigning computers or users to the correct role, which determines their network access privileges. When the IAS server receives a request for authentication, it compares the request with the list of remote access policies. The first policy to match the request is executed; additional policies are not searched.

Configuring Policies

The policies in this 802.1x authentication example are designed to work by examining the username portion of the authentication request, searching the Active Directory database for a matching name, and then examining the group membership for a computer or user entry that matches. For example, the following policies would operate with the controller configuration shown in Configuring Authentication with an 802.1X RADIUS Server on page 265:

- The Wireless-Computers policy matches the "Domain Computers" group. This group contains the list of all computers that are members of the domain. This group is used for all computers to authenticate to the network.
- The Wireless-Student policy matches the "Student" group. This group is used for all student users.
- The Wireless-Faculty policy matches the "Faculty" group. This group is used for all faculty users.
- The Wireless-Sysadmin policy matches the "Sysadmin" group. This group is used for system administrators.

In addition to matching the respective group, the policy also specifies that the request must be from an 802.11 wireless device. The policy instructs IAS to grant remote access permission if all the conditions specified in the policy match, a valid username/password is supplied, the user's or computer's remote access permission is set to "Allow".

To configure a policy:

1. In the **Internet Authentication Service** window, select **Remote Access Policies**.
2. To add a new policy, select **Action > New Remote Access Policy**. This launches a wizard that steps you through configuring the remote access policy.
3. Click **Next** on the initial wizard window to proceed.
4. Enter the name for the policy, for example, "Wireless Computers" and click **Next**.
5. In the **Access Method** window, select the **Wireless** option, then click **Next**.
6. in the **User or Group Access** window, select **Group** and click **Add** to add the group of users to which this policy applies (for example, "Domain Computers"). Click **Next**.
7. For Authentication Methods, select either **Protected EAP (PEAP) or Smart Card or other certificate**.
8. Click **Configure** to select additional properties.
9. Select a server certificate. The list of available certificates is taken from the computer certificate store on which IAS is running. In this case, a self-signed certificate was generated by the local certificate authority and installed on the IAS system. On each wireless client device, the local certificate authority is added as a trusted certificate authority, thus allowing this certificate to be trusted.

10. For PEAP, select the "inner" authentication method. The authentication method shown is MS-CHAPv2. (Because password authentication is being used on this network, this is the only EAP authentication type that should be selected.)

    You can also enable fast reconnect in this screen. If you enable fast reconnect here and also on client devices, additional time can be saved when multiple authentications take place (such as when clients are roaming between APs frequently) because the server will keep the PEAP encrypted tunnel alive.

11. Click **OK**.

Configuring RADIUS Attributes

In the configuration example for 802.1x, the controller restricts network access privileges based on the group membership of the computer or user. In order for this to work, the controller must be told to which group the user belongs. This is accomplished using RADIUS attributes returned by the authentication server.

To configure RADIUS attributes:

1. In the **Internet Authentication Service** window, select **Remote Access Policies**.
1. Open the remote access policy you want to configure, and select the **Advanced** tab.
2. Click **Add** to configure an attribute.
3. Select the **Class** attribute.
4. Enter the value for this attribute. For example, for the **Wireless-Computers** policy, the **Class** attribute returned to the controller should contain the value "computer".
5. Click **OK**.
6. Click **OK**.

Another example of a Class attribute configuration is shown below for the "Wireless-Student" policy. This policy returns the RADIUS attribute Class with the value "student" upon successful completion.

Configuring Management Authentication using IAS

Before you can configure the controller for management authentication using Windows IAS, you must perform the following steps to configure a Windows IAS RADIUS server on your Windows client.



The steps to perform this task may very depending on the version of Windows currently running on your server. For complete details on configuring Windows IAS, refer to the Windows documentation available at microsoft.com/downloads).

1. From your windows server, navigate to **Start > Settings > Control Panel > Administrative Tools>Internet Authentication Service**. The **Internet Authentication Service** window opens.
2. Verify that the Internet Authentication Service is running. If it is running, a green arrow icon will appear at the top of this window. If it has stopped, a red stop icon will appear. If the service is not active, click the green arrow icon to restart the service.
3. From the **Internet Authentication Service** window, right click the **Radius Clients** folder and select **New Radius Client**. The **New RADIUS Client** window opens.
4. Define a friendly name for the RADIUS client and enter the controller's IP address or DNS name. Click **Next**.
5. Enter and confirm the Shared Secret key for the controller then click **Finish**.

Next, create a remote policy for your new RADIUS client.

Creating a Remote Policy

1. From the **Internet Authentication Service** window, right click the **Remote Access Policies** folder and select **New Remote Access Policy**.
2. The **New Remote Access Policy** Wizard opens. Click **Next** on the first window to start the wizard.
3. Select **Use the wizard to set up a typical Policy for a common scenario** and enter a name for the policy, e.g Remote-Policy. Click **Next**.
4. In the **Access Method** window of the wizard, select the method you will use to gain management access to the network. Click **Next**.
5. In the **User or Group Access** window of the wizard, select either **user** or **group**, depending upon how your user permissions are defined. Click **Next**.

6. In the **Authentication Method** window, click the **Type** drop-down list and select **Protected EAP (PEAP)**. Click **Next**.

7. Click **Finish**.

Now you must define properties for the remote policy you just created.

Defining Properties for Remote Policy

1. In the **Internet Authentication Service** window, click the **Remote Access Policy** icon. All configured remote access policies will appear in the right window pane.
2. Right-click the policy you just created, and select **Properties**. The **Properties** window opens.
3. Select the **Grant remote access permission** radio button, and click **Edit Profile**. The **Edit Profile** window opens.
4. Click the **Authentication** tab and select the authentication methods that include **MS-CHAP, MS-CHAP V2** and **PAP**.
5. Click **Apply**.
6. Click the **Advanced** tab.
7. Click **Add**. The **Add Attribute** window opens.
8. Scroll down the list of attributes and select **Vendor-Specific**, then click **Add**. The **MultiValued Attribute Information** window appears.
9. Click **Add** again.
10. Enter the vendor code **14823** and select the option **Yes, It conforms**.
11. Click **Configure Attribute**. The **Configure VSA** window opens.
12. In the **Vendor-assigned attribute number** field, enter **3**.
13. In the **Attribute value** field, enter **7**.
14. Click **OK** to save your settings.
15. Click **Apply**.
16. Click **Apply**.

Now that you have defined your remote policy properties, you must create a user entry in the Windows active directory. The steps to complete this process will vary, depending on the version of Windows currently running on your server. The procedure below should be used only as a guideline.

Creating a User Entry in Windows Active Directory

1. Open the “Active Directory Users and Computers” tool on your Windows server.
2. Create a new user entry on the Windows Active directory.
3. Once you have created the new user, right-click the user name and select **Properties**.

4. Click the **Dial-in** tab and select **"Allow access"** for the user.
5. Click **Ok** to save your settings.

Configure the Controller to use IAS Management Authentication

The following procedure describes the steps to configure the controller to user IAS management authentication.

1. Access the controller WebUI and navigate to **Configuration>Authentication**.
2. Select the **Servers** tab.
3. Select **RADIUS** Server.
4. Enter a name for the RADIUS server in the entry field in the right window pane, then click **Add**.
5. Select the RADIUS server you just created from the list of servers in the left window pane to display configuration details for that server.
6. In the **Host** field, enter the IP address of the RADIUS server you want to use for Management Authentication.
7. Enter and then retype the shared key for the server.
8. Click **Apply**
9. Select **Server Group** from the server list on the left window pane.
10. In the entry blank on the right window pane, enter the name of a new server group (for example, "Management_group"), then click **Add**.
11. Click **Apply.**
12. Select the server group you just created from the list of server groups in the left window pane.
13. In the **Servers** section, click **New**.
14. Click the **Server Name** drop-down list and select your RADIUS server.
15. Click **Apply**.

Verify Communication between the Controller and the RADIUS Server

After you have configured your Windows Server and the Dell controller for Windows IAS Management Authentication, you can verify that the controller and server are communicating.

1. Navigate to **Diagnostics>AAA Test Server**.
2. Click the **Server Name** drop-down list and select the RADIUS server.
3. Select either **MSCHAP-V2** or **PAP** as the authentication method.
4. Enter the user name and password in the **Username** and **Password** fields.
5. Click **Begin Test**.
6. If the controller displays the words **Authentication Successful**, then the controller is able to communicate with the RADIUS server.

Window XP Wireless Client Sample Configuration

This section shows an example of how to configure a Windows XP wireless client using Windows XP's Wireless Zero Configuration service.



The following steps apply to a computer running Windows XP Professional Version 2002 with Service Pack 2. To configure a wireless client on other Windows platforms, see your Microsoft Windows documentation.

1. On the desktop, right-click My Network Places and select **Properties**.

2. In the Network Connections window, right-click on Wireless Network Connection and select **Properties**.
3. Select the **Wireless Networks** tab. This screen displays the available wireless networks and the list of preferred networks. Windows connects to the preferred networks in the order in which they appear in the list.
4. Click the **Advanced** button to display the Networks to access window.

   This window determines what types of wireless networks the client can access. By default, Windows connects to any type of wireless network. Make sure that the option Computer-to-computer (ad hoc) networks only is *not* selected. Click **Close**.
5. In the Wireless Networks tab, click **Add** to add a wireless network.
6. Click the **Association** tab to enter the network properties for the SSID.



This tab configures the authentication and encryption used between the wireless client and the Dell user-centric network. Therefore, the settings for the SSID that you configure on the client must *match* the configuration for the SSID on the controller.

- For an SSID using dynamic WEP, enter the following:
  - Network Authentication: Open
  - Data Encryption: WEP
  - Select the option "The key is provided for me automatically". Each client will use a dynamically-generated WEP key that is automatically derived during the 802.1x process.
- For an SSID using WPA, enter the following:
  - Network Authentication: WPA
  - Data Encryption: TKIP
- For an SSID using WPA-PSK, enter the following:
  - Network Authentication: WPA-PSK
  - Data Encryption: TKIP
  - Enter the pre-shared key.
- For an SSID using WPA2, enter the following:
  - Network Authentication: WPA2
  - Data Encryption: AES
- For an SSID using WPA2-PSK, enter the following:
  - Network Authentication: WPA2-PSK
  - Data Encryption: AES
  - Enter the pre-shared key



Do *not* select the option "This is a computer-to-computer (ad hoc) network; wireless access points are not used".

7. Click the **Authentication** tab to enter the 802.1x authentication parameters for the SSID. This tab configures the EAP type used between the wireless client and the authentication server.

   Configure the following, as shown in :
   - Select Enable IEEE 802.1x authentication for this network.
   - Select Protected EAP (PEAP) for the EAP type.
   - Select Authenticate as computer when computer information is available. The client will perform computer authentication when a user is not logged in.

- Do not select Authenticate as guest when user or computer information is unavailable. The client will not attempt to authenticate as a guest.
- Select Validate server certificate. This instructs the client to check the validity of the server certificate from an expiration, identity, and trust perspective.
- Select the trusted Certification Authority (CA) that can issue server certificates for the network.
- Select Secured password (EAP-MSCHAP v2) — the PEAP "inner authentication" mechanism will be an MS-CHAPv2 password.
- Select Enable Fast Reconnect to speed up authentication in some cases.

8. Under Select Authentication Method, click **Configure** to display the EAP-MSCHAPv2 Properties window. Select the option Automatically use my Windows logon name and password (and domain if any). This option specifies that the user's Windows logon information is used for authentication to the wireless network. This option allows the same logon credentials to be used for access to the Windows domain as well as the wireless network.

Acronyms

The following table lists the acronyms and their definitions used in this guide.

Table 235: *List of acronyms*

Acronym	Definition
ABR	area border router
AC	access category
ACI	adjacent channel interference
ACL	access control list
ADP	Aruba Discovery Protocol (ADP)
AES	advanced encryption standard
AIFSN	arbitrary inter-frame space number
ALG	application level gateway
AM	air monitor
AP	access point
APM	AP air monitor
ARM	adaptive radio management
AVF	AntiVirus Firewall
A-MSDU	aggregate MAC service data unit
BCMC	broadcast and multicast
BRAS	broadband remote access server
BRE	basic regular expression
BPDU	bridge protocol data unit
BSSID	basic service set identifier
CA	certification authority
CAC	call admission control
CAP	campus AP
CCA	clear channel assessment

Acronym	Definition
CDP	Cisco Discovery Protocol
CDR	call detail records
CHAP	Challenge Handshake Authentication Protocol
CRL	certificate revocation list
CSA	channel switch announcement
CSMA/CA	carrier sense multiple access with collision avoidance
CSR	certificate signing request
CTS	clear to send
CW	contention window
DAS	distributed antenna systems
DCF	distributed coordination function
DES	data encryption standard
DHCP	Dynamic Host Configuration Protocol
DS	differentiated services
DSCP	differentiated services codepoint
DSSS	direct sequence spread spectrum
DNS	domain name system
DoS	denial of service
DPD	dead peer detection
DR	designated router
DU	data unit
DMO	dynamic multicast optimization
EAP	Extensible Authentication Protocol
EAP-TLS	EAP-transport layer security
EDCA	enhanced distributed channel access
EIRP	effective isotropic radiated power
ESI	external service interfaces
ESS	extended service set

Acronym	Definition
ESSID	extended service set identifier
FE	fast ethernet
FFT	fast fourier transform
FHSS	frequency-hopping spread spectrum
FIB	forwarding information base
FRER	frame receive error rate
FRR	frame retry rate
FSPL	free space path loss
FTP	File Transfer Protocol
FQLN	fully qualified location name
GRE	generic routing encapsulation
GIS	generic interface specification
GMT	Greenwich Mean Time
GPP	guest provisioning page
HMD	high mobility device
HSPA	high-speed packet access
HT	high throughput
IAS	internet authentication server
IDS	intrusion detection system
IE	information element
IEEE	Institute of Electrical and Electronics Engineer
IGMP	Internet Group Management Protocol
IGP	Interior Gateway Routing Protocol
IKE PSK	internet key exchange pre-shared key
ISAKMP	Internet Security Association and Key Management Protocol
LACP	Link Aggregation Control Protocol
LAG	link aggregation group
LD	local debug

Acronym	Definition
LDAP	Lightweight Directory Access Protocol
LEAP	Lightweight Extensible Authentication Protocol
LI	listening interval
L2TP	Layer-2 Tunneling Protocol
MAC	media access control
MCS	modulation and coding scheme
MDPU	MAC protocol data unit
MIB	management information base
MIMO	multiple input, multiple output
MMS	mobility management system
MP	mesh point
MPP	mesh portal
MPV	mesh private VLAN
MSCHAP	Microsoft Challenge Handshake Authentication Protocol
MSCHAPv2	MSCHAP version 2
MSSID	mesh service set identifier
MPPE	Microsoft point-to-point encryption
MTU	maximum transmission unit
NAS	network access server
NAT	network address translation
NIC	network interface card
NOE	new office environment
NTP	Network Time Protocol
OCSP	Online Certificate Status Protocol
OFDM	orthogonal frequency division multiplexing
OKC	opportunistic key caching
OSPF	open shortest path first
OUI	organizationally unique identifier

Acronym	Definition
PAC	protected access credential
PAP	Password Authentication Protocol
PAPI	proprietary access protocol interface
PFS	perfect forward secrecy
PHB	per hop behavior
PIN	personal identification number
PKI	public key infrastructure
PMK	pairwise master key
PoE	power over ethernet
PSK	pre-shared key
PPPoE	point-to-point protocol over ethernet
PPTP	Point-to-Point Tunneling Protocol
PVST	per VLAN spanning tree
QoS	quality of service
RADIUS	remote authentication dial-in user service
RAP	remote AP
REGEX	region with the regular expression
RF	radio frequency
RFID	radio frequency identification
RoW	rest of world
RSSI	received signal strength indication
RSTP	Rapid Spanning Tree Protocol
RTLS	real-time locating systems
RTS	request to send
SA	security association
SDR	software-defined radio
SIM	subscriber identity module
SIP	Session Initiation Protocol

Acronym	Definition
SNIR	signal-to-noise-and-interference ratio
SNMP	Simple Network Management Protocol
SSID	service set identifier
STP	Spanning Tree Protocol
STRAP	secure thin remote access point
SVP	spectralink voice priority
TFTP	Trivial File Transfer Protocol
TIM	traffic indication map
TLS	transport layer security
TOS	type of service
TPM	trusted platform module
TSPEC	traffic specification
TXOP	opportunity to transmit
UDP	User Datagram Protocol
UTMS	universal mobile telecommunication systems
U-APSD	unscheduled automatic power save delivery
VBA	virtual branch networking
VIA	virtual intranet access
VoFI	voice over Wi-Fi
VoIP	voice over IP
VPN	virtual private network
VRD	validated reference design
VRRP	Virtual Router Redundancy Protocol
VSA	vendor specific attributes
VTP	Virtual Trunking Protocol
WIDS	wireless intrusion detection system
WINS	windows internet naming service
WIPS	wireless intrusion prevention system

Acronym	Definition
WISPr	wireless internet service provider roaming
WLAN	wireless local area network
WMM	wireless multimedia
WMS	WLAN management system
WSIRT	wireless security incident response team
WZC	wireless zero config
XAuth	extended authentication

Terms

The following table lists the terms and their definitions used in this guide.

Table 236: *List of terms*

Term	Definition
802.11	An evolving family of specifications for wireless LANs developed by a working group of the Institute of Electrical and Electronics Engineers (IEEE). 802.11 standards use the Ethernet protocol and CSMA/CA (carrier sense multiple access with collision avoidance) for path sharing.
802.11a	Provides specifications for wireless systems. Networks using 802.11a operate at radio frequencies in the 5GHz band. The specification uses a modulation scheme known as orthogonal frequency-division multiplexing (OFDM) that is especially well suited to use in office settings.
802.11b	WLAN standard often called Wi-Fi; backward compatible with 802.11. Instead of the phase-shift keying (PSK) modulation method historically used in 802.11 standards, 802.11b uses complementary code keying (CCK), which allows higher data speeds and is less susceptible to multipath-propagation interference.
802.11d	A wireless network communications specification for use in countries where systems using other standards in the 802.11 family are not allowed to operate. Configuration can be fine-tuned at the Media Access Control layer (MAC layer) level to comply with the rules of the country or district in which the network is to be used. Rules subject to variation include allowed frequencies, allowed power levels, and allowed signal bandwidth. 802.11d facilitates global roaming.
802.11e	A proposed adaptation to the 802.11a and 802.11b specifications that enhances the 802.11 Media Access Control layer (MAC layer) with a coordinated time division multiple access (TDMA) construct, and adds error-correcting mechanisms for delay-sensitive applications such as voice and video. The 802.11e specification provides seamless interoperability between business, home, and public environments such as airports and hotels and offers all subscribers high-speed Internet access with full-motion video, high-fidelity audio, and Voice over IP (VoIP).

Term	Definition
802.11g	Offers transmission over relatively short distances at up to 54 megabits per second (Mbps), compared with the 11 Mbps theoretical maximum of 802.11b. 802.11g employs orthogonal frequency division multiplexing (OFDM), the modulation scheme used in 802.11a, to obtain higher data speed. Computers or terminals set up for 802.11g can fall back to speeds of 11 Mbps, so that 802.11b and 802.11g devices can be compatible within a single network.
802.11h	Intended to resolve interference issues introduced by the use of 802.11a in some locations, particularly with military radar systems and medical devices. Dynamic frequency selection (DFS) detects the presence of other devices on a channel and automatically switches the network to another channel if and when such signals are detected. Transmit power control (TPC) reduces the radio-frequency (RF) output power of each network transmitter to a level that minimizes the risk of interference.
802.11i	Provides improved encryption for networks that use 802.11a, 802.11b, and 802.11g standards. Requires new encryption key protocols, known as Temporal Key Integrity Protocol (TKIP) and Advanced Encryption Standard (AES). Other features include key caching, which facilitates fast reconnection to the server for users who have temporarily gone offline, and pre-authentication, which allows fast roaming and is ideal for use with advanced applications such as Voice over Internet Protocol (VoIP).
802.11j	Proposed addition to the 802.11 family of standards that incorporates Japanese regulatory extensions to 802.11a; the main intent is to add channels in the radio-frequency (RF) band of 4.9 GHz to 5.0 GHz. WLANs using 802.11j will provide for speeds of up to 54 Mbps, and will employ orthogonal frequency division multiplexing (OFDM). The specification will define how Japanese 802.11 family WLANs and other wireless systems, particularly HiperLAN2 networks, can operate in geographic proximity without mutual interference.
802.11k	Proposed standard for how a WLAN should perform channel selection, roaming, and transmit power control (TPC) to optimize network performance. In a network conforming to 802.11k, if the access point (AP) having the strongest signal is loaded to capacity, a wireless device is connected to one of the under used APs. Even though the signal may be weaker, the overall throughput is greater because more efficient use is made of the network resources.
802.11n	Wireless networking standard to improve network throughput over the two previous standards 802.11a and 802.11g with a significant increase in the maximum raw data rate from 54 Mbit/s to 600 Mbit/s with the use of four spatial streams at a channel width of 40 MHz.
802.11m	An initiative to perform editorial maintenance, corrections, improvements, clarifications, and interpretations relevant to documentation for 802.11 family specifications. 802.11m also refers to the set of maintenance releases itself.
802.11 bSec	The bSec protocol is a pre-standard protocol that has been proposed to the IEEE 802.11 committee as an alternative to 802.11i. The difference between bSec and standard 802.11i is that bSec implements Suite B algorithms whenever possible. Notably, AES-CCM is replaced by AES-CGM, and the Key Derivation Function (KDF) of 802.11i is upgraded to

Term	Definition
	support SHA-256 and SHA-384. To provide interoperability with standard Wi-Fi software drivers, bSec is implemented as a shim layer between standard 802.11 Wi-Fi and a Layer 3 protocol such as IP. A controller configured to advertise a bSec SSID will advertise an open network, however only bSec frames will be permitted on the network.
802.1X	Standard designed to enhance 802.11 WLAN security. 802.1X provides an authentication framework, allowing a user to be authenticated by a central authority. The actual algorithm that is used to determine whether a user is authentic is left open and multiple algorithms are possible.
access point (AP)	An access point connects users to other users within the network and also can serve as the point of interconnection between the WLAN and a fixed wire network. The number of access points a WLAN needs is determined by the number of users and the size of the network.
access point mapping	The act of locating and possibly exploiting connections to WLANs while driving around a city or elsewhere. To do war driving, you need a vehicle, a computer (which can be a laptop), a wireless Ethernet card set to work in promiscuous mode, and some kind of an antenna which can be mounted on top of or positioned inside the car. Because a WLAN may have a range that extends beyond an office building, an outside user may be able to intrude into the network, obtain a free Internet connection, and possibly gain access to company records and other resources.
adhoc network	A LAN or other small network, especially one with wireless or temporary plug-in connections, in which some of the network devices are part of the network only for the duration of a communications session or, in the case of mobile or portable devices, while in some close proximity to the rest of the network.
A-MSDU	A structure containing multiple MSDUs , transported within a single (unfragmented) data medium access control (MAC) protocol data unit (MPDU).
band	A specified range of frequencies of electromagnetic radiation.
digital wireless pulse	Wireless technology for transmitting large amounts of digital data over a wide spectrum of frequency bands with very low power for a short distance. Ultra wideband radio can carry a huge amount of data over a distance up to 230 feet at very low power (less than 0.5 milliwatts), and has the ability to carry signals through doors and other obstacles that tend to reflect signals at more limited bandwidths and a higher power.
evil twin	A home-made wireless access point that masquerades as a legitimate one to gather personal or corporate information without the end-user's knowledge. It's fairly easy for an attacker to create an evil twin by simply using a laptop, a wireless card and some readily-available software. The attacker positions himself in the vicinity of a legitimate Wi-Fi access point and lets his computer discover what name and radio frequency the legitimate access point uses. He then sends out his own radio signal, using the same name.

Term	Definition
extensible authentication protocol (EAP)	Authentication protocol for wireless networks that expands on methods used by the point-to-point protocol (PPP), a protocol often used when connecting a computer to the Internet. EAP can support multiple authentication mechanisms, such as token cards, smart cards, certificates, one-time passwords, and public key encryption authentication.
fixed wireless	Wireless devices or systems in fixed locations such as homes and offices. Fixed wireless devices usually derive their electrical power from the utility mains, unlike mobile wireless or portable wireless which tend to be battery-powered. Although mobile and portable systems can be used in fixed locations, efficiency and bandwidth are compromised compared with fixed systems.
frequency allocation	Use of radio frequency spectrum regulated by governments.
frequency spectrum	Part of the electromagnetic spectrum.
goodput	Goodput is the ratio of the total bytes transmitted or received in the network to the total air time required for transmitting or receiving the bytes. The air time includes the retransmission time taken for both successful and dropped frames.Suppose 1000 frames of 1500 bytes each are transmitted in the network as follows: • 50% of frames are transmitted successfully at MCS index 11 at 108 Mbps. • 25% of the frames were dropped in the 1st attempt at 108 Mbps but were successfully transmitted using MCS index 3 at 54 Mbps in the second attempt. • The remaining 25% are dropped in both the attempts. Then the effective rate is calculated as: The total bits transmitted / the total air time. In this example: (500 * 1500 + 250 * 1500) * 8 / (total air time for 50% frames + total air time for 25 % frames retransmitted + total air time for 25% dropped frames) = 40.5 Mbps.
hot spot	A WLAN node that provides Internet connection and virtual private network (VPN) access from a given location. A business traveller, for example, with a laptop equipped for Wi-Fi can look up a local hot spot, contact it, and get connected through its network to reach the Internet and their own company remotely with a secure connection. Increasingly, public places, such as airports, hotels, and coffee shops are providing free wireless access for customers.
hot zone	A wireless access area created by multiple hot spots located in close proximity to each other. Hot zones usually combine public safety access points with public hot spots. Each hot spot typically provides network access for distances between 100 and 300 feet; various technologies, such as mesh network topologies and fiber optic backbones, are used in conjunction with the hot spots to create areas of coverage.
Infrared Data Association(IrDA)	An industry-sponsored organization set up in 1993 to create international standards for the hardware and software used in infrared communication links. In this special form of radio transmission, a focused ray of light in the infrared frequency spectrum, measured in terahertz, or trillions of hertz (cycles per second), is modulated with information and sent from a transmitter to a receiver over a relatively short distance

Term	Definition
IR wireless	The use of wireless technology in devices or systems that convey data through infrared (IR) radiation. Infrared is electromagnetic energy at a wavelength or wavelengths somewhat longer than those of red light. The shortest-wavelength IR borders visible red in the electromagnetic radiation spectrum; the longest-wavelength IR borders radio waves.
microwave	Electromagnetic energy having a frequency higher than 1 gigahertz (billions of cycles per second), corresponding to wavelength shorter than 30 centimeters. Microwave signals propagate in straight lines and are affected very little by the troposphere. They are not refracted or reflected by ionized regions in the upper atmosphere. Microwave beams do not readily diffract around barriers such as hills, mountains, and large human-made structures.
MIMO	An antenna technology for wireless communications in which multiple antennas are used at both the source (transmitter) and the destination (receiver). The antennas at each end of the communications circuit are combined to minimize errors and optimize data speed. MIMO is one of several forms of smart antenna technology, the others being MISO (multiple input, single output) and SIMO (single input, multiple output).
MISO	An antenna technology for wireless communications in which multiple antennas are used at the source (transmitter). The antennas are combined to minimize errors and optimize data speed. The destination (receiver) has only one antenna. MISO is one of several forms of smart antenna technology, the others being MIMO (multiple input, multiple output) and SIMO (single input, multiple output).
near field communication(NFC)	A short-range wireless connectivity standard (Ecma-340, ISO/IEC 18092) that uses magnetic field induction to enable communication between devices when they're touched together, or brought within a few centimeters of each other. The standard specifies a way for the devices to establish a peer-to-peer (P2P) network to exchange data.
optical wireless	The combined use of conventional radio-frequency (RF) wireless and optical fiber for telecommunication. Long-range links are provided by optical fiber and links from the long-range end-points to end users are accomplished by RF wireless or laser systems. RF wireless at ultra-high frequencies (UHF) and microwave frequencies can carry broadband signals to individual computers at substantial data speeds.
OCSP Client	The ArubaOScontroller can act as an OCSP client and issues OCSP queries to remote OCSP responders located on the intranet or Internet.
OCSP Responder	The OCSP client retrieves certificate revocation status from an OCSP responder. The responder may be the certificate authority (CA) that has issued the certificate in question or it may be some other designated entity which provides the service on behalf of the CA.
radio frequency (RF)	Portion of electromagnetic spectrum in which electromagnetic waves are generated by feeding alternating current to an antenna.
structured wireless-aware network (SWAN)	A technology that incorporates a WLAN into a wired wide-area network (WAN). SWAN technology can enable an existing wired network to serve hundreds of users, organizations, corporations, or agencies over a large geographic area. A SWAN is said to be scalable, secure, and reliable.
secure copy (SCP)	Secured encrypted command to copy files across an ssh connection,

Term	Definition
	Files can be copied from or to a remote server, and also from one remote server to another.
transponder	A wireless communications, monitoring, or control device that picks up and automatically responds to an incoming signal. The term is a contraction of the words transmitter and responder. Transponders can be either passive or active.
ultra high frequency (UHF)	International Telecommunication Union (ITU) band 9, 300-3000 MHz, 1m - 100 mm frequency wavelength.
ultra wideband (UVB)	Is a wireless technology for transmitting large amounts of digital data over a wide spectrum of frequency bands with very low power for a short distance. Ultra wideband broadcasts very precisely timed digital pulses on a carrier signal across a very wide spectrum (number of frequency channels) at the same time. UWB can carry a huge amount of data over a distance up to 230 feet at very low power (less than 0.5 milliwatts), and has the ability to carry signals through doors and other obstacles that tend to reflect signals at more limited bandwidths and a higher power.
virtual private network (VPN)	A network that uses a public telecommunication infrastructure, such as the Internet, to provide remote offices or individual users with secure access to their organization's network. A VPN ensures privacy through security procedures and tunneling protocols such as the Layer Two Tunneling Protocol (L2TP). Data is encrypted at the sending end and decrypted at the receiving end.
voice over WLAN (VoWLAN)	A method of routing telephone calls for mobile users over the Internet using the technology specified in IEEE 802.11b. Routing mobile calls over the Internet makes them free, or at least much less expensive than they would be otherwise.
wideband code-division multiple access (W-CDMA)	Officially known as IMT-2000 direct spread; ITU standard derived from Code-Division Multiple Access (CDMA). W-CDMA is a third-generation (3G) mobile wireless technology that promises much higher data speeds to mobile and portable wireless devices than commonly offered in today's market.
Wi-Fi	A term for certain types of WLANs. Wi-Fi can apply to products that use any 802.11 standard. Wi-Fi has gained acceptance in many businesses, agencies, schools, and homes as an alternative to a wired LAN. Many airports, hotels, and fast-food facilities offer public access to Wi-Fi networks.
WiMAX	A wireless industry coalition whose members organized to advance IEEE 802.16 standards for broadband wireless access (BWA) networks. WiMAX 802.16 technology is expected to enable multimedia applications with wireless connection and, with a range of up to 30 miles, enable networks to have a wireless last mile solution. According to the WiMAX forum, the group's aim is to promote and certify compatibility and interoperability of devices based on the 802.16 specification, and to develop such devices for the marketplace.

Term	Definition
wired equivalent privacy (WEP)	A security protocol specified in 802.11b, designed to provide a WLAN with a level of security and privacy comparable to what is usually expected of a wired LAN. Data encryption protects the vulnerable wireless link between clients and access points; once this measure has been taken, other typical LAN security mechanisms such as password protection, end-to-end encryption, virtual private networks (VPNs), and authentication can be put in place to ensure privacy.
wireless	Describes telecommunications in which electromagnetic waves (rather than some form of wire) carry the signal over part or all of the communication path.
wireless abstract XML (WAX)	Describes telecommunications in which electromagnetic waves (rather than some form of wire) carry the signal over part or all of the communication path.
wireless application service provider (WASP)	Provides Web-based access to applications and services that would otherwise have to be stored locally and makes it possible for customers to access the service from a variety of wireless devices, such as a smartphone or personal digital assistant (PDA).
wireless ISP (WISP)	An internet service provider (ISP) that allows subscribers to connect to a server at designated hot spots (access points) using a wireless connection such as Wi-Fi. This type of ISP offers broadband service and allows subscriber computers, called stations, to access the Internet and the Web from anywhere within the zone of coverage provided by the server antenna, usually a region with a radius of several kilometers.
wireless service provider	A company that offers transmission services to users of wireless devices through radio frequency (RF) signals rather than through end-to-end wire communication.
wireless local area network (WLAN)	A local area network (LAN) that users access through a wireless connection. 802.11 standards specify WLAN technologies. WLANs are frequently some portion of a wired LAN.
yagi antenna	A unidirectional antenna commonly used in communications when a frequency is above 10 MHz.