

Dell PowerConnect

ArubaOS 5.0

User Guide



Copyright

© 2010 Aruba Networks, Inc. AirWave®, Aruba Networks®, Aruba Mobility Management System®, and other registered marks are trademarks of Aruba Networks, Inc. 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. Any other trademarks appearing in this manual are the property of their respective companies.

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. The Open Source code used can be found at this site:

http://www.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

	About this Guide.....	37
	Audience.....	37
	Fundamentals.....	37
	WebUI.....	37
	CLI.....	37
	Related Documents.....	38
	Conventions.....	38
	Contacting Support.....	39
Chapter 1	The Basic User-Centric Networks.....	41
	Configuring the User-Centric Network.....	41
	Deployment and Configuration Tasks.....	41
	Deployment Scenario #1.....	42
	Deployment Scenario #2.....	42
	Deployment Scenario #3.....	43
	Configuring the Controller.....	44
	Running the Initial Setup.....	44
	Connecting to the Controller after Initial Setup.....	45
	Configuring a VLAN for Network Connection.....	45
	Creating and Updating a VLAN.....	46
	Viewing Existing VLAN IDs.....	46
	Creating, Updating, and Deleting VLAN Pools.....	46
	Adding existing VLAN IDs to a VLAN Pool in the CLI.....	46
	Assigning and Configuring the Trunk Port.....	47
	In the WebUI.....	47
	In the CLI.....	47
	Configuring the Default Gateway.....	47
	In the WebUI.....	47
	In the CLI.....	47
	Configuring the Loopback for the Controller.....	47
	In the WebUI.....	48
	In the CLI.....	48
	Configuring the System Clock.....	48
	Installing Licenses.....	49
	Connecting the Controller to the Network.....	49
	Deploying APs.....	49
	Running the RF Plan.....	49
	Enabling APs to Connect to the Controller.....	50
	Enabling APs to Obtain IP Addresses.....	50
	Enabling the DHCP server on the controller in the WebUI.....	50
	Enable the DHCP server on the controller in the CLI.....	50
	Locating the Controller.....	51
	From a DNS Server.....	51
	From a DHCP Server.....	51
	Using the Dell Discovery Protocol (ADP).....	52
	Provisioning APs for Mesh.....	52
	Installing APs.....	52
	Updating the RF Plan.....	53

	Additional Configuration.....	53
Chapter 2	Network Parameters.....	55
	Configuring VLANs	55
	Creating and Updating VLANs	55
	Using the WebUI	55
	Using CLI.....	55
	Create a Bulk VLANs using the WebUI	56
	Using CLI.....	56
	Creating, Updating and Deleting VLAN Pools	56
	Creating a VLAN pool using the WebUI	56
	Updating a VLAN Pool	56
	Deleting a VLAN Pool	56
	Create a VLAN Pool using CLI.....	57
	Viewing existing VLAN IDs using CLI.....	57
	Adding existing VLAN IDs using CLI	57
	Configuring Ports.....	57
	Classifying Traffic as Trusted or Untrusted	58
	About Trusted and Untrusted Physical Ports	58
	About Trusted and Untrusted VLANs.....	58
	Configuring Trusted/Untrusted Ports and VLANs.....	58
	Using WebUI	58
	Using CLI.....	59
	Configure Trusted/Untrusted Ports and VLANs in Trunk Mode	59
	Using the WebUI	59
	Using CLI.....	60
	About VLAN Assignments.....	60
	Assigning a Static Address to a VLAN	60
	Using the WebUI	60
	Using CLI.....	60
	Configuring a VLAN to Receive a Dynamic Address.....	60
	Enabling the DHCP Client.....	61
	Using the WebUI	61
	Using the CLI	61
	Enabling the PPPoE Client	61
	Using the WebUI	62
	Using CLI.....	62
	Default Gateway from DHCP/PPPoE	62
	Using the WebUI	62
	Using CLI.....	62
	Configuring DNS/WINS Server from DHCP/PPPoE.....	62
	Using the WebUI.....	62
	Using CLI.....	63
	Configuring Source NAT to Dynamic VLAN Address.....	63
	Using the WebUI	63
	Using CLI.....	63
	Configuring Source NAT for VLAN Interfaces	63
	Example Configuration	64
	Using the WebUI	64
	Using CLI.....	64
	Inter-VLAN Routing.....	64
	Using the WebUI to restrict VLAN routing.....	65
	Using CLI.....	65
	Configuring Static Routes	65
	Using the WebUI.....	65
	Using CLI	66
	Configuring the Loopback IP Address	66

Using the WebUI.....	66
Using CLI	66
Using the CLI to reboot the controller.....	66
Configuring the Controller IP Address	66
Using CLI	67
Configuring GRE Tunnels.....	67
Creating a Tunnel Interface	68
Using the WebUI	68
Using CLI.....	68
Directing Traffic into the Tunnel	68
Static Routes	68
Firewall Policy.....	68
Tunnel Keepalives	69

Chapter 3

RF Plan.....	71
Supported Planning.....	71
Before You Begin.....	72
Task Overview.....	72
Planning Requirements	72
Launching the RF Plan	74
Campus List Page	74
Building List Pane.....	75
Building Specifications Overview	75
Building Dimension Page	76
AP Modeling Parameters Page	77
Radio Type.....	78
Design Model.....	78
Overlap Factor	79
Users/AP	79
Radio Properties (Desired Rates and HT Support Options)	80
AM Modeling Page	81
Design Models.....	82
Monitor Rates	82
Planning Floors Page	82
Zoom.....	83
Approximate Coverage Map	83
Floor Editor Dialog Box.....	84
Area Editor Dialog Box	85
Access Point Editor Dialog Box	86
AP Plan Page.....	89
Initialize.....	89
Optimize	89
Fix All Suggested AP/AMs.....	89
AM Plan Page	90
Initialize.....	90
Optimize	90
Fix All Suggested AP/AMs.....	90
Exporting and Importing Files.....	90
Export Campus.....	91
Import Campus.....	91
Export Buildings Page	91
Import Buildings Page	92
Locate.....	92
FQLN Mapper	92
Using the FQLN Mapper in the AP Provision Page.....	94
Using the WebUI.....	94
Using CLI	95

RF Plan Example	95
Sample Building.....	95
Create a Building	96
Model the Access Points	97
Model the Air Monitors	97
Add and Edit a Floor.....	97
Adding the background image and naming the first floor.....	98
Adding the background image and naming the second floor.....	98
Defining Areas	98
Creating a Don't Care Area	98
Creating a Don't Deploy Area.....	99
Running the AP Plan	99
Running the AM Plan	100

Chapter 4

Access Points.....	101
Remote AP vs Campus AP	101
Basic Configuration.....	102
AP Names and Groups	102
AP Names.....	102
Duplicate Names.....	103
Renaming in the WebUI	103
Renaming in the CLI	103
AP Groups.....	103
Creating an AP group in the WebUI.....	104
Assigning APs to an AP group in the WebUI.....	104
Creating an AP group in the CLI.....	104
Assigning an AP to an AP group in the CLI.....	104
Virtual APs	105
Configuring Profiles.....	105
Wireless LAN Profiles.....	106
AP Profiles	107
QoS Profiles.....	108
RF Management Profiles.....	108
IDS Profiles.....	109
Mesh Profiles	109
Controller Profiles.....	110
Profile Hierarchy.....	110
Applying Profiles.....	113
Excluding a virtual AP profile from an AP in the WebUI.....	115
Excluding a virtual AP profile from an AP in the CLI	115
Viewing Profile Errors.....	115
Virtual AP Configurations	115
Configuring the WLAN.....	116
Configuring the User Role	116
In the WebUI	117
In the CLI	117
Configuring Authentication Servers.....	117
In the WebUI	117
In the CLI	117
Configuring Authentication.....	118
In the WebUI	118
In the CLI	119
Applying the Virtual AP.....	119
In the WebUI	119
Creating a new SSID Profile.....	119
In the CLI	123

Guest WLAN.....	123
Configuring the VLAN	123
In the WebUI	123
In the CLI	123
Configuring the Guest Role	123
In the WebUI	123
In the CLI	124
Configuring the Guest Virtual AP	124
In the WebUI	124
In the CLI	125
Configuring High-throughput on Virtual APs	125
In the WebUI	125
In the CLI	127
Managing High-throughput Radio Profiles	127
Advanced Configuration Options.....	127
Configuring 802.11k Protocol	127
In the WebUI	128
In the CLI	129
Configuring RF Optimization	129
In the WebUI	129
In the CLI	130
Configuring RF Event Configuration	130
In the WebUI	130
In the CLI	131
Changing AP Installation Modes	132
In the WebUI	132
In the CLI	132
Configuring Channel Switch Announcement (CSA)	133
In the WebUI	133
In the CLI	133
20 MHz and 40 MHz Static Channel Assignments	133
Configuring Channels in the WebUI	134
Configuring channels in the CLI	135
Automatic Channel and Transmit Power Selection Using ARM	135
APs Over Low-Speed Links.....	135
Best Practices.....	135
Configuring in the WebUI.....	136
Configuring in the CLI	136
Prioritizing AP heartbeats in the WebUI	136
Prioritizing AP heartbeats in the CLI	136
AP Redundancy	136
AP failback.....	136
Configuring in the WebUI.....	137
Configuring in the CLI	137
AP Maintenance Mode	137
Configuring in the WebUI.....	137
Configuring in the CLI	137
Managing AP LEDs.....	138
Disabling LEDs in the WebUI.....	138
Enable or Disable LEDs in the CLI.....	138
Configuring Blinking LEDs in the CLI.....	138
Chapter 5 Adaptive Radio Management (ARM).....	139
ARM Overview	139
ARM Support for 802.11n.....	139
Monitoring Your Network with ARM.....	140

Noise and Error Monitoring	140
Application Awareness	140
ARM Profiles	140
Creating a New ARM Profile	141
Copying an Existing Profile	141
Deleting a Profile	142
Configuring ARM Settings via the WebUI	142
Configuring ARM Settings via the CLI	145
Assigning an ARM Profile to an AP Group	146
Assigning a Profile via the WebUI	146
Assigning a Profile via the CLI	146
Multi-Band ARM and 802.11a/802.11g Traffic	147
Band Steering	147
Enabling Band Steering via the WebUI	147
Enabling Band Steering via the CLI	148
Traffic Shaping	148
Enabling Traffic Shaping via the WebUI	148
Spectrum Load Balancing	149
RX Sensitivity Tuning Based Channel Reuse	149
Non-802.11 Noise Interference Immunity	150
ARM Metrics	150
ARM Troubleshooting	151
Too many APs on the Same Channel	151
Wireless Clients Report a Low Signal Level	152
Transmission Power Levels Change Too Often	152
APs Detect Errors but Do Not Change Channels	152
APs Don't Change Channels Due to Channel Noise	152

Chapter 6

Remote Access Points	153
Overview	153
Configuring the Secure Remote Access Point Service	154
Configure a Public IP Address for the Controller	155
Using the WebUI to create a DMZ address	155
Using CLI	155
Configure the VPN Server	155
Using the WebUI	155
Using CLI	156
Configure the Remote AP User Role	156
Using the WebUI	156
Using CLI	157
Configure VPN Authentication	158
Using the WebUI	158
Using CLI	158
Configuring Internal Database for Authentication	158
Using the WebUI	159
Configure VPN authentication using the internal database	160
Add the user to the internal database	160
Using CLI to configure the internal DB for a RAP user	160
Provision the AP	161
Creating a Remote AP Whitelist	162
Revoking an AP	162
Deploying a Branch Office/Home Office Solution	162
Configuring the branch office AP	163
Troubleshooting Remote AP	163
Local Debugging	163

Remote AP Summary.....	164
Remote AP Connectivity.....	166
Remote AP Diagnostics.....	166
Enabling Double Encryption.....	166
Using the WebUI	167
Using CLI.....	167
Advanced Configuration Options.....	167
Understanding Remote AP Modes of Operation.....	167
Fallback Mode.....	169
Configuring the fallback mode	169
Using WebUI to configure the AAA profile.....	170
Using CLI.....	170
Using the WebUI to configure virtual AP profile.....	170
Using CLI.....	171
Configuring the DHCP Server on the Remote AP.....	171
Using the WebUI	172
Using CLI.....	173
Advanced Backup Configuration Options.....	173
Using the WebUI to configure the session ACL.....	174
Using the WebUI to configure the AAA profile	175
Using the WebUI to define the backup configuration.....	175
Using the CLI to configure the session ACL.....	176
Using the CLI to configure the AAA profile.....	176
Using the CLI to define the backup configuration	176
DNS Controller Setting	177
Specify the DNS name using the WebUI	177
Backup Controller List	177
Configuring the LMS and backup LMS IP addresses using WebUI	178
Configuring the LMS and backup LMS IP addresses using CLI.....	178
Remote AP Failback.....	178
Using the WebUI	178
Using the CLI	179
RAP Local Network Access	179
Using the WebUI	179
Using CLI.....	179
Remote AP Authorization Profiles	180
Add or Edit a Remote AP Authorization Profile	180
Access Control Lists and Firewall Policies.....	180
Split Tunneling	180
Configuring Split Tunneling	181
Configuring the Session ACL.....	182
Using the WebUI	182
Using the CLI	183
Configuring ACL for restricted LD homepage access.....	183
Using CLI.....	184
Configuring the AAA Profile and the Virtual AP Profile	184
Using the WebUI	184
Using CLI.....	185
Configuring split tunneling in the virtual AP profile.....	185
Using the CLI to configure split tunneling in the virtual AP profile.....	185
Using the WebUI to list the corporate DNS servers	186
Using the CLI to list the corporate DNS servers	186
Wi-Fi Multimedia	186
Uplink Bandwidth Reservation.....	186
Bandwidth Reservation for Uplink Voice Traffic.....	187
Configuring Bandwidth Reservation	187
Using the WebUI	187

	Using CLI	187
Chapter 7	Secure Enterprise Mesh.....	189
	Mesh Access Points	189
	Mesh Portals	190
	Mesh Points.....	190
	Mesh Clusters	190
	Mesh Links.....	191
	Link Metrics.....	192
	Optimizing Links	192
	Mesh Profiles	192
	Mesh Cluster Profile	193
	Mesh Radio Profile.....	193
	RF Management (802.11a and 802.11g) Profiles	193
	Adaptive Radio Management Profiles.....	193
	High-Throughput Profiles.....	194
	Mesh High-Throughput SSID Profile	194
	Wired AP Profile	194
	Mesh Recovery Profile	194
	Mesh Solutions	195
	Thin AP Services with Wireless Backhaul Deployment.....	195
	Point-to-Point Deployment	196
	Point-to-Multipoint Deployment	196
	High-Availability Deployment.....	197
	Before You Begin.....	198
	Pre-Deployment Considerations.....	198
	Outdoor-Specific Deployment Considerations.....	198
	Configuration Considerations.....	198
	Post-Deployment Considerations	198
	AP-12x Specific Considerations	199
	Mesh Radio Profiles.....	199
	Managing Mesh Profiles via the WebUI	199
	Creating a New Profile	199
	Assigning a Profile to a Mesh AP or AP Group.....	202
	Editing a Profile.....	202
	Deleting a Profile.....	203
	Managing Mesh Profiles using the CLI	203
	Creating or Modifying a Profile	203
	Viewing Profile Settings.....	203
	Assigning a Profile to an AP Group.....	203
	Deleting a Mesh Radio Profile	204
	RF Management (802.11a and 802.11g) Profiles	204
	Managing 802.11a/802.11g Profiles Via the WebUI	204
	Creating a Profile.....	204
	Assigning an 802.11a/802.11g Profile	207
	Assigning a High-throughput Profile.....	207
	Assigning an ARM Profile.....	208
	Editing an 802.11a/802.11g Profile.....	209
	Deleting a Profile.....	209
	Managing 802.11a/802.11g Profiles using the CLI	209
	Creating or Modifying a Profile	209
	Viewing RF Management Settings.....	210
	Assigning a 802.11a/802.11g Profile	210
	Deleting a Profile.....	211
	Mesh High-Throughput SSID Profiles.....	211
	Managing Profiles via the WebUI	211

Creating a Profile.....	211
Assigning a Profile to an AP Group.....	212
Editing a Profile.....	212
Deleting a Profile.....	213
Managing high-throughput SSID profiles using the CLI.....	213
Creating or Modifying a Profile.....	213
Assigning a Profile to an AP Group.....	213
Viewing High-throughput SSID Settings.....	213
Deleting a Profile.....	214
Mesh Cluster Profiles.....	214
Deployments with Multiple Mesh Cluster Profiles.....	214
Managing Mesh Cluster Profiles via the WebUI.....	215
Creating a Profile.....	215
Associating a Profile to Mesh APs.....	216
Editing a Profile.....	217
Deleting a Mesh Cluster Profile.....	217
Managing Mesh Cluster Profiles Using the CLI.....	217
Viewing Mesh Cluster Profile Settings.....	218
Associating Mesh Cluster Profiles.....	218
Excluding a Mesh Cluster Profile from a Mesh Node.....	219
Deleting a Mesh Cluster Profile.....	219
Ethernet Ports for Mesh.....	219
Configure bridging on the Ethernet port.....	219
Ethernet Ports for Secure Jack Operation.....	220
Configuring Secure Jack via the WebUI.....	220
Configuring Secure Jack via the CLI.....	221
Extending the Life of a Mesh Network.....	221
Modifying the AP System Profile via the WebUI.....	221
Modifying the AP System Profile via the CLI.....	221
Provisioning Mesh Nodes.....	222
Outdoor AP Parameters.....	222
Provisioning Caveats.....	223
Provisioning Mesh Nodes via the WebUI.....	223
Provisioning Mesh Nodes via the CLI.....	224
AP Boot Sequence.....	224
Mesh Portal.....	224
Mesh Point.....	224
Air Monitoring and Mesh.....	224
Verifying the Network.....	225
Remote Mesh Portals.....	225
How RMP Works.....	225
Creating a Remote Mesh Portal via the WebUI.....	226
Provisioning the AP.....	226
Defining the Mesh Private VLAN.....	227
Selecting a Mesh Radio Profile.....	228
Selecting an RF Management Profile.....	228
Adding a Mesh Cluster Profile.....	228
Configuring a DHCP Pool.....	229
Configuring the VLAN ID of the Virtual AP Profile.....	229
Provisioning a Remote Mesh Portal via the CLI.....	230
Additional Information.....	230

Chapter 8	Authentication Servers.....	231
	Important Points to Remember.....	231
	Servers and Server Groups.....	231
	Configuring Servers.....	232

Configuring a RADIUS Server	232
In the WebUI	233
In the CLI	233
RADIUS Server Authentication Codes.....	233
Configuring an LDAP Server.....	234
In the WebUI	234
In the CLI	234
Configuring a TACACS+ Server	235
In the WebUI	235
In the CLI	235
Configuring a Windows Server.....	235
In the WebUI	236
In the CLI	236
Internal Database.....	236
Configuring the Internal Database	236
In the WebUI	237
In the CLI	237
RAP Static Inner IP Address	237
In the WebUI	237
In the CLI	238
Managing Internal Database Files	238
Exporting files in the WebUI.....	238
Importing files in the WebUI.....	239
In the CLI	239
Internal Database Utilities	239
Deleting All User.....	239
Repairing the Internal Database.....	239
Server Groups	239
Configuring Server Groups	239
In the WebUI	240
In the CLI	240
Configuring Server List Order and Fail-Through	240
In the WebUI	240
In the CLI	241
Configuring Dynamic Server Selection	241
In the WebUI	242
In the CLI	242
Configuring Match FQDN Option	243
In the WebUI	243
In the CLI	243
Trimming Domain Information from Requests.....	243
In the WebUI	243
In the CLI	244
Configuring Server-Derivation Rules	244
In the WebUI	245
In the CLI	245
Configuring a Role Derivation Rule for the Internal Database	245
In the WebUI	245
In the CLI	246
Assigning Server Groups	246
User Authentication	246
Management Authentication	246
In the WebUI	246
In the CLI	246
Accounting	247
RADIUS Accounting.....	247
In the WebUI	248

In the CLI	249
TACACS+ Accounting.....	249
Configuring Authentication Timers.....	249
Setting an Authentication Timer	249
In the WebUI	250
In the CLI	250

Chapter 9

802.1x Authentication	251
Overview of 802.1x Authentication	251
Supported EAP Types	252
Authentication with a RADIUS Server	252
Authentication Terminated on Controller.....	253
Configuring 802.1x Authentication	254
Using the WebUI	255
Using the CLI	259
Configuring and Using Certificates with AAA FastConnect.....	260
Using the WebUI	260
Using the CLI	260
Configuring User and Machine Authentication	260
Role Assignment with Machine Authentication Enabled.....	261
VLAN Assignment with Machine Authentication Enabled.....	262
Example Configurations.....	262
Authentication with an 802.1x RADIUS Server.....	262
Configuring Roles and Policies	263
Creating the student role and policy	263
Creating the faculty role and policy	264
Creating the guest role and policy	265
Creating roles and policies for sysadmin and computer	266
Creating an alias for the internal network using CLI.....	267
Configuring the RADIUS Authentication Server	267
Using the WebUI	267
Using the CLI	268
Configure 802.1x Authentication	268
Using the WebUI	268
Using the CLI	268
Configure VLANs	269
Using the WebUI	269
Using the CLI	269
Configuring the WLANs.....	270
Configuring the Guest WLAN	270
Using the WebUI	270
Using the CLI	271
Configuring the Non-Guest WLANs	271
Using the WebUI	271
Using the CLI	272
Authentication with the Controller's Internal Database	272
Configuring the Internal Database	272
Using the WebUI	272
Using the CLI	273
Configuring a server rule using the WebUI	273
Configuring a server rule using the CLI	273
Configure 802.1x Authentication	273
Using the WebUI	273
Using the CLI	274
Configure VLANs	274
Using the WebUI	274
Using the CLI	275

Configuring the WLANs.....	275
Configuring the Guest WLAN	275
Using the WebUI	275
Using the CLI	276
Configuring the Non-Guest WLANs	276
Using the WebUI	276
Using the CLI	277
Advanced Configuration Options for 802.1x.....	278
Configuring reauthentication with Unicast Key Rotation	278
Using the WebUI	278
Using the CLI	278

Chapter 10

Roles and Policies	279
Policies.....	279
Access Control Lists (ACLs).....	280
Creating a Firewall Policy	280
In the WebUI	282
In the CLI	282
Creating an ACL White List.....	282
Configuring a White List Bandwidth Contract in the WebUI	282
Configuring the ACL White List in the WebUI.....	282
Configuring the White List Bandwidth Contract in the CLI.....	283
Configuring the ACL White List in the CLI	283
User Roles.....	283
Creating a User Role	284
In the WebUI	284
In the CLI	285
Bandwidth Contracts	285
Configuring a Bandwidth Contract in the WebUI	285
Assigning a Bandwidth Contract to a User Role in the WebUI	285
Configuring and Assigning Bandwidth Contracts in the CLI.....	286
User Role Assignments	286
User Role in AAA Profile	286
In the WebUI	286
In the CLI	287
User-Derived Role	287
Configuring a User-derived Role	288
Configure a User-derived Role in the CLI.....	288
Default Role for Authentication Method	288
In the WebUI	288
In the CLI	289
Server-Derived Role.....	289
VSA-Derived Role.....	289
Global Firewall Parameters	289

Chapter 11

Stateful and WISPr Authentication	293
Stateful Authentication Overview.....	293
WISPr Authentication Overview.....	293
Important Points to Remember	294
Stateful 802.1x Authentication.....	294
Configure Authentication via the WebUI	294
Configure Authentication via the CLI	295
Stateful NTLM Authentication.....	295
Configure Authentication via the WebUI	295
Configure Authentication via the CLI	296

Configuring WISPr Authentication	296
Configure WISPr Authentication via the WebUI.....	296
Configure WISPr Authentication via the CLI	297

Chapter 12

Captive Portal	299
Captive Portal Overview	299
Policy Enforcement Firewall Next Generation (PEFNG) License	299
Controller Server Certificate	300
Captive Portal in the Base ArubaOS	300
Configuring Captive Portal via the WebUI	301
Configuring Captive Portal via the CLI	302
Captive Portal with the PEFNG License.....	302
Configuring Captive Portal via the WebUI	303
Configuring Captive Portal via the CLI	304
Example Authentication with Captive Portal	305
Creating a Guest-logon User Role	305
Creating an Auth-guest User Role.....	306
Configure Policies and Roles via the WebUI	306
Time Range.....	306
Auth-Guest-Access Policy	307
Block-Internal-Access Policy	308
Drop-and-Log Policy	309
Guest-logon Role.....	309
Guest-Logon Role.....	309
Configure Policies and Roles via the CLI.....	310
Time Range.....	310
Create Aliases.....	310
Guest-Logon-Access Policy.....	310
Auth-Guest-Access Policy	310
Block-Internal-Access Policy	311
Drop-and-Log Policy	311
Guest-Logon Role.....	311
Auth-Guest Role	311
Guest VLANs	311
Configuring the guest VLAN via the WebUI.....	311
Configuring the guest VLAN via the CLI	312
Captive Portal Authentication	312
Modifying the Initial User Role.....	313
Configuring the AAA Profile	313
Configuring the WLAN.....	313
User Account Administration	314
Captive Portal Configuration Parameters	314
Optional Captive Portal Configurations.....	316
Per-SSID Captive Portal Page.....	316
Changing the Protocol to HTTP.....	317
Proxy Server Redirect.....	318
Redirecting Clients on Different VLANs	319
Web Client Configuration with Proxy Script	319
Personalizing the Captive Portal Page.....	320

Chapter 13

Advanced Security	323
Securing Client Traffic	324
Securing Wireless Clients.....	324
In the WebUI	325
In the CLI	325
Securing Wired Clients.....	325

In the WebUI	326
In the CLI	327
Securing Wireless Clients Through Non-Dell APs	327
In the WebUI	327
In the CLI	328
Securing Controller-to-Controller Communication	328
Configuring Controllers for xSec	329
In the WebUI	329
In the CLI	329
Configuring the Odyssey Client on Client Machines	329
Installing the Odyssey Client	329

Chapter 14 Virtual Intranet Access333

VIA.....	333
How it Works.....	333
Installing VIA.....	334
Upgrade Workflow.....	334
Configuring the VIA Controller	334
Before you Begin.....	334
Configuring VIA Settings	335
Using WebUI to Configure VIA.....	335
Enable VPN Server Module	335
Create VIA User Roles.....	335
Create VIA Authentication Profile	336
Create VIA Connection Profile	337
Configure VIA Web Authentication.....	340
Associate VIA Connection Profile to User Role	340
Configure VIA Client WLAN Profiles	341
Re-branding VIA and Downloading the Installer	343
Using CLI to Configure VIA	344
Enable VPN module	344
Create VIA roles	344
Create VIA authentication profiles.....	345
Create VIA connection profiles.....	345
Configure VIA web authentication	345
Associate VIA connection profile to user role	345
Configure VIA client WLAN profiles.....	345
Customize VIA logo, landing page and downloading installer	345
VPN Configuration.....	347

Chapter 15 Virtual Private Networks347

VPN authentication.....	348
Supported VPN AAA Deployments	348
Remote Access VPN for L2TP IPsec	349
Configure the VPN via the WebUI	349
Authentication Method and Server Addresses	349
Define Address Pools	349
Source NAT.....	349
IKE Shared Secrets.....	350
IKE Policies.....	350
Configure the VPN via the CLI	350
Authentication Method and Server Addresses	350
Address Pools.....	350
Source NAT.....	350
IKE Shared Secrets.....	350
IKE Policies.....	351

Example Configurations for Remote Access Clients.....	351
Configuring a VPN for Smart Card Clients via the WebUI.....	351
Configuring a VPN for Smart Card Clients via the CLI.....	353
VPNs for L2TP/IPsec Clients with Passwords.....	353
Configure the L2TP/IPsec VPN via the WebUI.....	354
Configure the L2TP/IPsec VPN via the CLI.....	355
Remote Access VPNs for XAuth.....	355
Configure a VPN with XAuth via the WebUI.....	355
Authentication Method and Server Addresses	356
Address Pools.....	356
Source NAT.....	356
Aggressive Mode	356
Server Certificate	356
CA Certificate for VPN Clients.....	356
IKE Shared Secrets.....	357
IKE Policies.....	357
Configure a VPN with XAuth via the CLI.....	357
Authentication Method and Server Addresses	357
Address Pools.....	357
Source NAT.....	357
Aggressive Mode	357
Server Certificate	357
CA Certificate Assigned for VPN Clients.....	357
IKE Shared Secrets.....	357
IKE Policies.....	358
Configurations for XAuth Clients using Smart Cards	358
Configure a VPN for Cisco XAuth Smart Card Clients.....	358
Configure the VPN via the WebUI	358
Configure the VPN via the CLI.....	359
XAuth Clients Using a Username/Password	360
Configure the VPN via the WebUI	360
Configure the VPN via the CLI.....	361
Remote Access VPN for PPTP	362
Configuring a VPN with PPTP via the WebUI.....	362
Configuring a VPN with PPTP via the CLI.....	362
Site-to-Site VPNs.....	362
Site-to-Site VPNs with Dynamic IP Addresses.....	363
VPN Topologies	363
Configuring site-to-site VPN via the WebUI	363
Configuring site-to-site VPN via the CLI	365
Dead Peer Detection.....	366
Dell Dialer	366
Configuring the Dialer via the WebUI	366
Configuring the Dialer via the CLI	367
Captive Portal Dialer	367
Configuring MAC-Based Authentication.....	369
Configuring the MAC Authentication Profile	369
Chapter 16 MAC-based Authentication	369
Using the WebUI to configure a MAC authentication profile.....	370
Using the CLI to configure a MAC authentication profile.....	370
Configuring Clients	370
Using the WebUI to configure clients in the internal database	370
Using the CLI to configure clients in the internal database	371
Chapter 17 Control Plane Security.....	373

Control Plane Security Overview	373
Configuring Control Plane Security	374
Verifying Certificates	374
Configuring Settings via the WebUI	375
Configuring Settings via the CLI	376
Managing the Campus AP Whitelist	376
Viewing Entries in the Campus AP Whitelist	377
Modifying an AP in the Campus AP Whitelist	378
Revoking an AP via the Campus AP Whitelist.....	379
Deleting an AP Entry from the Campus AP Whitelist.....	380
Purging the Campus AP Whitelist	380
Whitelists on Master and Local Controllers	380
Campus AP Whitelist Synchronization	381
Viewing and Managing the Master or Local controller Whitelists.....	382
Viewing the Master or Local Controller Whitelist	382
Deleting an Entry from the Master or Local Controller Whitelist.....	382
Purging the Master or Local Controller Whitelist.....	383
Environments with Multiple Master Controllers	383
Configuring Networks with a Backup Master Controller	383
Configuring Networks with Clusters of Master Controllers	384
Creating a Cluster Root	384
Creating a Cluster Member	385
Viewing Controller Cluster Settings	385
Replacing a Controller on a Multi-Controller Network	386
Replacing Controllers in a Single Master Network.....	386
Replacing a Local Controller	386
Replacing a Master Controller (With No Backup).....	387
Replacing a Redundant Master Controller	387
Replacing Controllers in a Multi-Master Network	388
Replacing a Local Controller in a Multi-Master Network	388
Replacing a Cluster Member Controller (With no Backup).....	388
Replacing a Redundant Cluster Member Controller	389
Replacing a Cluster Root Controller with no Backup Controller	389
Replacing a Redundant Cluster Root Controller	389
Troubleshooting Control Plane Security.....	390
Certificate Problems	390
Disabling Control Plane Security.....	390
Verify Whitelist Synchronization	390
Supported APs	391
Rogue APs	391

Chapter 18	Adding Local Controllers	393
	Moving to a Multi-Controller Environment.....	393
	Preshared Key for Inter-Controller Communication.....	393
	Best Security Practices for the Preshared Key	394
	Configuring the Preshared Key	394
	Using the WebUI to configure the Local Controller PSK.....	394
	Using the WebUI to configure the Master Controller PSK.....	394
	Using the CLI to configure the PSK	395
	Configuring Local Controllers.....	395
	Configuring the Local Controller	395
	Using the Initial Setup	395
	Using the Web UI	396
	Using the CLI	396
	Configuring Layer-2/Layer-3 Settings	396
	Configuring Trusted Ports	396

	Configuring APs	396
	Using the WebUI to configure the LMS IP	396
	Using the CLI to configure the LMS IP	397
Chapter 19	IP Mobility	399
	Dell Mobility Architecture	399
	Configuring Mobility Domains	400
	Configuring a Mobility Domain	401
	Using the WebUI	401
	Using the CLI	401
	Joining a Mobility Domain	402
	In the WebUI	402
	In the CLI	402
	Example Configuration.....	402
	Configuring Mobility using the WebUI.....	403
	Configuring Mobility using the CLI	404
	Tracking Mobile Users	404
	Mobile Client Roaming Status	404
	Viewing mobile client status using the WebUI	404
	Viewing mobile client status using the CLI	404
	Viewing user roaming status using the CLI	405
	Viewing specific client information using the CLI	405
	Mobile Client Roaming Locations	405
	In the WebUI	406
	In the CLI	406
	HA Discovery on Association	406
	Setting up mobility association using CLI	406
	Advanced Mobility Functions	406
	Configuring advanced mobility functions using the WebUI.....	406
	Configuring mobility functions using CLI	408
	Proxy Mobile IP	408
	Proxy DHCP	409
	Revocations.....	409
	Mobility Multicast.....	409
	Proxy IGMP and Proxy Remote Subscription	409
	Inter-controller Mobility	410
	Configuring Mobility Multicast Using the WebUI	411
	Configuring Mobility Multicast Using the CLI	411
	Example.....	411
Chapter 20	VRRP	413
	Redundancy Parameters.....	413
	Configuring the Local Controller for Redundancy	414
	In the WebUI	414
	In the CLI	415
	Configuring the LMS IP	415
	In the WebUI	415
	In the CLI	415
	Configuring the Master Controller for Redundancy	415
	Configuring Database Synchronization	417
	In the WebUI	417
	In the CLI	417
	Configuring Master-Local Controller Redundancy.....	418
Chapter 21	RSTP.....	421
	Migration and Interoperability	421

Rapid Convergence.....	421
Edge Port and Point-to-Point.....	422
Configuring RSTP.....	422
In the WebUI	422
In the CLI	424
Monitoring RSTP.....	424
Troubleshooting.....	424

Chapter 22

W-600 Series Controller.....	427
Important Points to Remember	427
Internal Access Point (AP).....	428
USB Cellular Modems.....	428
Functional Description.....	428
Mode-Switching	428
USB Modems Commands	428
Uplink Manager	429
Cellular Profile.....	430
Dialer Group	431
Configuring a Supported USB Modem.....	432
Configuring a New USB Modem	433
Configuring the Profile and Modem Driver	433
Configuring the TTY Port	435
Testing the TTY Port.....	436
Selecting the Dialer Profile.....	436
Linux Support	437
NAS (Network-Attached Storage).....	437
NAS Device Setup.....	437
Configuring in the CLI.....	438
Managing NAS Devices	439
Mounting and Unmounting Devices	439
In the WebUI	440
Print Server.....	442
Printer Setup	442
In the CLI	442
In the WebUI	443
Other commands for managing printers	443
Sample Topology and Configuration	444
Remote Branch 1—W-651 Controller	444
Remote Branch 2—W-650 Controller	445
W-3200 Central Office Controller—Active	446
W-3200 Central Office Controller—Backup.....	448

Chapter 23

OSPFv2.....	451
Important Points to Remember	451
WLAN Scenario	451
WLAN Topology.....	452
WLAN Routing Table.....	452
Branch Office Scenario.....	453
Branch Office Topology.....	453
Branch Office Routing Table	454
Configuring OSPF.....	455
Deployment Best Practices	456
Sample Topology and Configuration	457
Remote Branch 1	457

Remote Branch 2.....	458
W-3200 Central Office Controller—Active.....	459
W-3200 Central Office Controller—Backup.....	461

Chapter 24 Wireless Intrusion Prevention..... 463

IDS Features.....	463
Unauthorized Device Detection.....	463
Rogue/Interfering AP Detection.....	463
Adhoc Network Detection and Containment.....	464
Wireless Bridge Detection.....	464
Misconfigured AP Detection.....	464
Weak WEP Detection.....	464
Multi Tenancy Protection.....	464
MAC OUI Checking.....	464
Denial of Service (DoS) Detection.....	464
Rate Analysis.....	465
Fake AP.....	465
Impersonation Detection.....	465
Station Disconnection.....	465
EAP Handshake Analysis.....	465
Sequence Number Analysis.....	465
AP Impersonation.....	465
Signature Detection.....	466
IDS Configuration.....	466
IDS Profile Hierarchy.....	466
Configuring IDS via the WebUI.....	466
Configuring IDS via the CLI.....	467
IDS General Profile.....	467
Configuring the General Profile via the WebUI.....	468
Configuring the General Profile via the CLI.....	468
Denial of Service Attack Detection.....	468
Configuring the DoS profile via the WebUI.....	470
Configuring the DoS profile via the CLI.....	470
IDS Rate Thresholds Profile.....	471
Configuring the Rate Threshold Profile via the WebUI.....	471
Configuring the Rate Thresholds Profile via the CLI.....	472
Impersonation Detection Profile.....	472
Configuring the Impersonation Profile via the WebUI.....	472
Configuring the Impersonation Profile via the CLI.....	473
Signature Matching Profile.....	473
Configuring the Signature Matching Profile via the WebUI.....	473
Configuring the Signature Matching Profile via the CLI.....	474
IDS Signature Profile.....	474
Creating a New Signature via the WebUI.....	474
Creating a New Signature via the CLI.....	475
Unauthorized Device Detection.....	475
Configuring the Unauthorized Device Profile via the WebUI.....	479
Configuring the Unauthorized Device Profile via the CLI.....	479
WLAN Management System.....	480
Configuring WMS via the WebUI.....	480
Configuring WMS via the CLI.....	481
Configuring Local WMS Settings.....	481
Managing the WMS Database.....	481
Enabling AP Learning.....	481
Classifying APs.....	482
Configuring Misconfigured AP Detection and Protection.....	482
Valid Enterprise SSIDs.....	483

	Client Blacklisting	484
	Methods of Blacklisting.....	484
	Manual Blacklisting	485
	Authentication Failure Blacklisting	485
	Attack Blacklisting	485
	Blacklist Duration	486
	Removing a Client from Blacklisting	487
Chapter 25	Link Aggregation Control Protocol.....	489
	Important Points to Remember	489
	Configuring LACP.....	489
	In the CLI	489
	In the WebUI	491
	Best Practices.....	491
	Sample Configuration	492
Chapter 26	Management Access.....	493
	Certificate Authentication for WebUI Access.....	493
	Configuring Certificate Authentication for WebUI Access.....	493
	In the WebUI	493
	In the CLI	494
	Public Key Authentication for SSH Access	494
	In the WebUI	494
	In the CLI	495
	Radius Server Authentication	495
	Radius Server Username/Password Authentication	495
	In the WebUI	495
	In the CLI	495
	RADIUS Server Authentication with VSA	496
	RADIUS Server Authentication with Server-Derivation Rule	496
	Configuring a Value-of Server-derivation Rule in the WebUI.....	496
	In the CLI	497
	Configuring a set-value server-derivation rule in the WebUI.....	497
	In the CLI	498
	Disabling Authentication of Local Management User Accounts.....	498
	In the WebUI	498
	In the CLI	498
	Verifying the configuration	498
	Resetting the Admin or Enable Password.....	498
	Setting an Administrator Session Timeout	499
	Setting a CLI Session Timeout	499
	Setting a WebUI Session Timeout.....	500
	Management Password Policy	500
	Defining a Management Password Policy.....	500
	In the WebUI	500
	Managed RFprotect Sensors.....	502
	Setting RFprotect Sensor Mode in the Radio Profile	502
	In the WebUI	502
	In the CLI	503
	Configuring the RFprotect Server	503
	In the WebUI	503
	In the CLI	503
	Reverting Managed Sensors to APs	503
	Managing Certificates	503
	About Digital Certificates	504

Obtaining a Server Certificate	504
In the WebUI	504
In the CLI	505
Obtaining a Client Certificate	505
Importing Certificates	505
In the WebUI	506
In the CLI	506
Viewing Certificate Information	506
Imported Certificate Locations	507
Checking CRLs	507
Configuring SNMP	507
SNMP Parameters for the Controller	507
In the WebUI	508
In the CLI	508
Configuring Logging	509
In the WebUI	510
In the CLI	510
Guest Provisioning	510
Configuring the Guest Provisioning Page	511
In the WebUI	511
Configuring the SMTP Server and Port in the WebUI.....	514
Configuring an SMTP server and port in the CLI	514
Creating Email Messages in the WebUI.....	514
Configuring a Guest Provisioning User	515
In the WebUI	516
In the CLI	517
Customizing the Guest Access Pass.....	517
Creating Guest Accounts	518
Guest Provisioning User Tasks	518
Optional Configurations.....	520
Restricting one Captive Portal Session for each Guest.....	520
Setting the Maximum Time for Guest Accounts	521
Managing Files on the Controller.....	521
Transferring ArubaOS Image Files	522
In the WebUI	522
In the CLI	522
Backing Up and Restoring the Flash File System	522
Backup the flash file system in the WebUI.....	522
Backup the flash file system in the CLI.....	523
Restore the flash file system in the WebUI.....	523
Restore the flash file system using CLI.....	523
Copying Log Files.....	523
In the WebUI	523
In the CLI	523
Copying Other Files	523
In the WebUI	524
In the CLI	524
Setting the System Clock	524
Manually Setting the Clock.....	524
In the WebUI	524
In the CLI	524
Configuring an NTP Server	524
In the WebUI	525
In the CLI	525
Chapter 27	
Software Licenses	527
Terminology	527

Licenses	528
License Types	528
Multi-Controller Network	529
License Usage.....	529
Interaction	530
Best Practices.....	530
Installing a License	531
Enabling a new license on your controller	531
Software License Email.....	531
Locating the System Serial Number	531
Obtaining a Software License Key	532
Creating a software license key	532
Applying the Software License Key in the WebUI.....	532
Applying the Software License Key in the License Wizard	532
Deleting a License	532
Moving Licenses.....	533
Resetting the Controller.....	533

Chapter 28

IPv6 Client Support	535
About IPv6.....	535
Support for IPv6	535
Supported Network Configuration	535
Network Connection for Windows IPv6 Clients	536
Features that Support IPv6	537
Authentication.....	537
Firewall Functions	537
Firewall Policies.....	539
Creating an IPv6 firewall policy	540
Assigning an IPv6 Policy to a User Role.....	541
DHCPv6 Passthrough/Relay	541
Multicast Snooping	541
IPv6 User Addresses.....	542
Viewing or Deleting User Entries.....	542
Viewing Datapath Statistics for IPv6 Sessions	542
Important Points to Remember	542

Chapter 29

Voice and Video	545
License Requirements	545
Configuring Voice	545
Setting up Net Services	545
Using Default Net Services	545
Creating or Modifying Net Services.....	546
Configuring User Roles.....	546
Using the Default User Role	546
Creating or Modifying User Roles	547
Using User-Derivation Roles	548
Optional Configurations.....	549
WPA Fast Handover.....	549
Mobile IP Home Agent Assignment	550
The VoIP Call Admission Control Profile	550
VoIP-Aware ARM Scanning.....	552
Voice-Aware 802.1x	552
SIP Authentication Tracking	552
Dial Plan for SIP Calls.....	553

Dial Plan Format	553
Configuring Dial Plans	554
Voice over Remote Access Point	556
Configuring Video	557
Configuring Video over WLAN enhancements	557
Pre-requisites.....	557
Using CLI	558
Using WebUI	560
QoS.....	561
Wi-Fi Multimedia	561
Using the WebUI to enable WMM	562
Using the CLI to enable WMM	562
Configurable WMM AC Mapping	562
Mapping Considerations	563
Using the WebUI to map between WMM AC and DSCP	563
Using the CLI to map between WMM AC and DSCP	564
Battery Boost	564
Using the WebUI to enable battery boost.....	564
Using the CLI to enable battery boost.....	564
Dynamic WMM Queue Management.....	565
Enhanced Distributed Channel Access.....	565
Using the WebUI to configure EDCA parameters.....	566
Using the CLI to configure EDCA parameters	567
WMM Queue Content Enforcement.....	567
Using the WebUI to enable WMM queue content enforcement.....	567
Using the CLI to enable WMM queue content enforcement.....	567

Chapter 30	External Services Interface	569
	Understanding ESI.....	569
	Understanding the ESI Syslog Parser.....	571
	ESI Parser Domains	571
	Peer Controllers.....	573
	Syslog Parser Rules	573
	Condition Pattern Matching	573
	User Pattern Matching.....	574
	ESI Configuration Overview.....	574
	Configuring Health-Check Method, Groups, and Servers.....	575
	In the WebUI	575
	In the CLI	575
	Defining the ESI Server	575
	In the WebUI	576
	In the CLI	576
	Defining the ESI Server Group	576
	In the WebUI	576
	In the CLI	577
	Redirection Policies and User Role.....	577
	In the WebUI	577
	In the CLI	578
	ESI Syslog Parser Domains and Rules	578
	Managing Syslog Parser Domains in the WebUI	578
	Adding a new syslog parser domain.....	578
	Deleting an existing syslog parser domain.....	579
	Editing an existing syslog parser domain.....	579
	Managing Syslog Parser Domains in the CLI	579
	Adding a new syslog parser domain.....	579
	Showing ESI syslog parser domain information	579
	Deleting an existing syslog parser domain.....	579

Editing an existing syslog parser domain.....	579
Managing Syslog Parser Rules.....	580
In the WebUI	580
Adding a new parser rule	580
Deleting a syslog parser rule	581
Editing an existing syslog parser rule.....	581
Testing a Parser Rule	581
In the CLI	582
Adding a new parser rule	582
Showing ESI syslog parser rule information:	582
Deleting a syslog parser rule:	582
Editing an existing syslog parser rule.....	582
Testing a parser rule.....	582
Monitoring Syslog Parser Statistics	582
In the WebUI	582
In the CLI	583
Example Route-mode ESI Topology	583
ESI server configuration on controller	583
IP routing configuration on Fortinet gateway.....	583
Configuring the Example Routed ESI Topology.....	583
Health-Check Method, Groups, and Servers.....	584
Defining the Ping Health-Check Method.....	584
In the WebUI	584
In the CLI	584
Defining the ESI Server	585
In the WebUI	585
In the CLI	585
Defining the ESI Server Group	585
In the WebUI	585
In the CLI	586
Redirection Policies and User Role.....	586
In the WebUI	586
In the CLI	587
Syslog Parser Domain and Rules	587
Add a New Syslog Parser Domain in the WebUI.....	587
Adding a New Parser Rule in the WebUI.....	588
In the CLI	588
Example NAT-mode ESI Topology.....	588
ESI server configuration on the controller.....	589
Configuring the Example NAT-mode ESI Topology	590
Configuring the NAT-mode ESI Example in the WebUI	590
In the WebUI	590
Configuring the ESI Group in the WebUI.....	591
Configure the ESI Servers in the WebUI	591
Configuring the Redirection Filter in the WebUI	591
Configuring the Example NAT-mode Topology in the CLI	592
Configuring a Health-Check Ping	592
Configuring ESI Servers	592
Configure an ESI Group, Add the Health-Check Ping and ESI Servers.....	592
Using the ESI Group in a Session Access Control List	592
CLI Configuration Example 1.....	592
CLI Configuration Example 2.....	593
Basic Regular Expression Syntax.....	593
Character-Matching Operators	594
Regular Expression Repetition Operators	594
Regular Expression Anchors.....	595
References	595

Appendix A	DHCP with Vendor-Specific Options	597
	Overview	597
	Windows-Based DHCP Server.....	597
	Configuring Option 60.....	597
	To configure option 60 on the Windows DHCP server.....	598
	Configuring Option 43.....	598
	To configure option 43 on the Windows DHCP server:.....	598
	Linux DHCP Servers	599
Appendix B	External Firewall Configuration	601
	Communication Between Dell Devices	601
	Network Management Access	602
	Other Communications	602
Appendix C	Behavior and Defaults	603
	Mode Support	603
	Basic System Defaults.....	604
	Network Services.....	604
	Policies.....	606
	Roles	608
	Default Management User Roles.....	610
	Default Open Ports	613
Appendix D	802.1x Configuration for IAS and Windows Clients.....	617
	Configuring Microsoft IAS	617
	RADIUS Client Configuration	617
	Remote Access Policies.....	618
	Active Directory Database.....	618
	Configuring Policies.....	618
	Configuring RADIUS Attributes.....	620
	Configure Management Authentication using IAS.....	621
	Configure the Controller to use IAS Management Authentication.....	622
	Verify Communication between the Controller and the RADIUS Server	624
	Window XP Wireless Client Example Configuration	624
Appendix E	Internal Captive Portal	629
	Creating a New Internal Web Page	629
	Basic HTML Example.....	630
	Installing a New Captive Portal Page	631
	Displaying Authentication Error Message	631
	Reverting to the Default Captive Portal	632
	Language Customization	632
	Customizing the Welcome Page	635
	Customizing the Pop-Up box.....	636
	Customizing the Logged Out Box	637
Appendix F	VIA End User Instructions	639
	Pre-requisites.....	639
	Downloading VIA.....	639
	Installing VIA	640
	Using VIA	640

	Connection Details Tab	640
	Diagnostic Tab	641
	Diagnostics Tools.....	641
	Settings Tab.....	641
Appendix G	Provisioning RAP at Home	643
	Provision the RAP using a Static IP Address	643
	Provision the RAP on a PPPoE Connection	644
	Using 3G/EVDO USB Modem.....	645
Index		649

Figures

Figure 1	APs Connected to Controller	53
Figure 2	IP Address Assignment to VLAN via DHCP or PPPoE.....	61
Figure 3	Example: Source NAT using Controller IP Address.....	64
Figure 4	Default Inter-VLAN Routing	65
Figure 5	Plan>Campus List Window	74
Figure 6	Plan>Building List Pane.....	75
Figure 7	Plan>New Building>Overview Window	75
Figure 8	Plan>New Building>Specification Window.....	76
Figure 9	Plan>New Building>AP Modeling Parameters Window	77
Figure 10	AM Modeling Page	81
Figure 11	Planning Floors Page	82
Figure 12	Coverage Map Example	84
Figure 13	Floor Editor Dialog Box	84
Figure 14	Area Editor Dialog Box	85
Figure 15	Access Point Editor	87
Figure 16	AP Planning	89
Figure 17	AP Groups	104
Figure 18	Virtual AP Configurations Applied to the same AP.....	105
Figure 19	AP Specific and AP Group Profile Hierarchies	111
Figure 20	Layer 2/Layer3 Profile Hierarchies	112
Figure 21	Excluding a Virtual AP Profile from an AP.....	114
Figure 22	Profile Errors	115
Figure 23	Remote AP with a Private Network.....	153
Figure 24	Remote AP with Controller on Public Network	154
Figure 25	Remote AP with Controller Behind Firewall.....	154
Figure 26	Remote AP in a Multi-Controller Environment.....	154
Figure 27	Remote AP with Single Controller	162
Figure 28	Sample Backup Controller Scenario.....	178
Figure 29	Enable Remote AP Local Network Access	179
Figure 30	Sample Split Tunnel Environment.....	181
Figure 31	Enable Restricted Access to LD Homepage	184
Figure 32	Uplink Bandwidth Reservation.....	187
Figure 33	Sample Mesh Clusters.....	191
Figure 34	Sample Wireless Backhaul Deployment.....	196
Figure 35	Sample Point-to-Point Deployment.....	196
Figure 36	Sample Point-to-Multipoint Deployment	197
Figure 37	Sample High-Availability Deployment.....	197
Figure 38	Working of RMP.....	226
Figure 39	Provisioning an AP as a Remote Mesh Portal.....	227
Figure 40	Server Group	232
Figure 41	IP-Address parameter in the local database.....	237
Figure 42	IP-Address parameter in the RAP Whitelist	238
Figure 43	Domain-Based Server Selection Example	242
Figure 44	802.1x Authentication with RADIUS Server	253

Figure 45	802.1x Authentication with Termination on Controller	254
Figure 46	Wireless xSec Client Example.....	324
Figure 47	Wired xSec Client Example.....	326
Figure 48	Controller-to-Controller xSec Example	329
Figure 49	The regedit Window.....	330
Figure 50	Modifying a regedit Policy	330
Figure 51	The Funk Odyssey Client Profile.....	331
Figure 52	Certificate Information.....	331
Figure 53	Network Profile.....	332
Figure 54	VIA - Associate User Role to VIA Authentication Profile	336
Figure 55	VIA - Creating a new server group for VIA authentication profile	336
Figure 56	VIA - Enter a name for the server group.....	337
Figure 57	VIA - Create VIA Connection Profile---redo image.....	337
Figure 58	VIA - Configure VIA Connection Profile.....	338
Figure 59	VIA - Select VIA Authentication Profile.....	340
Figure 60	VIA - Associate VIA Connection Profile to User Role	341
Figure 61	VIA - Create VIA Client WLAN Profile.....	341
Figure 62	VIA - Configure the SSID Profile	342
Figure 63	VIA - Configure VIA Client WLAN Profile	342
Figure 64	VIA - Customize VIA logo, Landing Page, and download VIA Installer	343
Figure 65	Site-to-Site VPN Configuration Components.....	363
Figure 66	Control Plane Security Settings	376
Figure 67	Local Controller Whitelist on a Master Controller	381
Figure 68	A Cluster of Master Controllers using Control Plane Security	384
Figure 69	Sequence numbers on Master and Local Controllers	391
Figure 70	Routing of Traffic to Mobile Client within Mobility Domain.....	400
Figure 71	Example Configuration: Campus-Wide	403
Figure 72	Redundant Topology: Master-Local Redundancy	418
Figure 73	Configuring RSTP.....	423
Figure 74	Monitoring RSTP.....	424
Figure 75	Cellular Profile Commands	429
Figure 76	Uplink Commands.....	429
Figure 77	Connected Cellular Devices	429
Figure 78	WebUI Uplink Manager.....	430
Figure 79	Cellular Profile from the WebUI	431
Figure 80	Configuring Dialer Group.....	432
Figure 81	Display supported USB modems	432
Figure 82	show usb verbose example (partial).....	432
Figure 83	show uplink.....	433
Figure 84	uplink cellular priority.....	433
Figure 85	show usb command	433
Figure 86	show usb verbose for profile and driver.....	434
Figure 87	cellular profile new_card command	434
Figure 88	Driver options.....	434
Figure 89	Driver=(none)	435
Figure 90	show usb ports 13 command.....	435
Figure 91	show usb test command	435
Figure 92	Time out error example.	435
Figure 93	Port I/O error	436
Figure 94	Device Ready State.....	436

Figure 95	usb test extended.	436
Figure 96	show dialer group example	437
Figure 97	W-650 Sample Topology.....	444
Figure 98	WLAN OSPF Topology	452
Figure 99	Branch Office OSPF Topology.....	453
Figure 100	General OSPF Configuration	455
Figure 101	Edit OSPF VLAN Settings.....	455
Figure 102	Monitoring OSPF.....	456
Figure 103	Sample OSPF Topology	457
Figure 104	Resetting the Password	499
Figure 105	Reconfigure the enable mode password	499
Figure 106	Guest Provisioning Configuration Page—Guest Fields Tab.....	512
Figure 107	Guest Provisioning Configuration Page—Page Design Tab.....	513
Figure 108	Guest Provisioning Configuration Page—Email Tab.....	515
Figure 109	Sample Guest Account Email – Sent to Sponsor	515
Figure 110	Customized Guest Account Information Window.....	517
Figure 111	Creating a Guest Account—Management User Summary Page.....	518
Figure 112	Creating a Guest Account—New Guest Window	519
Figure 113	Creating a Guest Account—Show Details Pop-up Window.....	519
Figure 114	Printing Guest Account Information.....	520
Figure 115	Supported Network Configuration.....	536
Figure 116	Setting DSCP value	560
Figure 117	Enabling Dynamic Multicast Optimization for Video and Set Threshold.....	560
Figure 118	Enabling Video Aware Scan	561
Figure 119	Configuring bandwidth management.....	561
Figure 120	ESI-Fortinet Topology	570
Figure 121	Load Balancing Groups	571
Figure 122	ESI Parser Domains	572
Figure 123	Peer Controllers.....	573
Figure 124	Example Route-Mode Topology.....	583
Figure 125	Example NAT-Mode Topology	589
Figure 126	Scope Options Dialog Box.	599
Figure 127	DHCP Scope Values.....	599
Figure 128	IAS RADIUS Clients.....	617
Figure 129	IAS Remote Access Policies	619
Figure 130	Policy Configuration Wizard—Authentication Methods	619
Figure 131	Policy Configuration Wizard—PEAP Properties	620
Figure 132	RADIUS class Attribute Configuration	620
Figure 133	Example RADIUS Class Attribute for “student”	621
Figure 134	Configuring a RADIUS Server for IAS Management Authentication.....	623
Figure 135	Configuring a Server Group for IAS Management Authentication	623
Figure 136	Testing a RADIUS Server	624
Figure 137	Wireless Networks.....	624
Figure 138	Networks to Access.....	625
Figure 139	Wireless Network Association	626
Figure 140	Wireless Network Authentication	627
Figure 141	Protected EAP Properties	627
Figure 142	EAP MSCHAPv2 Properties	628
Figure 143	Default Welcome Page.....	635
Figure 144	Login to Download VIA	639

Figure 145	Downloading VIA set up file after authentication.....	640
Figure 146	Show Advanced Settings.....	643
Figure 147	Provision RAP using Static IP.....	644
Figure 148	Provision RAP on a PPPoE Connection	645
Figure 149	Provision using a pre-configured USB Modem	646
Figure 150	Provision using a USB Modem with Custom Settings.....	646

Tables

Table 1	Typographical Conventions	38
Table 2	Alcatel-Lucent Contacts.....	39
Table 3	Web site contact	39
Table 4	Classifying Trusted and Untrusted Traffic.....	58
Table 5	Planning Worksheet - Building Dimensions	73
Table 6	Planning Worksheet - AP Desired Rates.....	73
Table 7	Planning Worksheet - AM Desired Rates	73
Table 8	Definition of Campus List Buttons	74
Table 9	Building List Buttons	75
Table 10	New Building Specifications Parameters	76
Table 11	AP Modeling Parameters.....	78
Table 12	Radio Type Definitions.....	78
Table 13	Design Model Radio Buttons.....	79
Table 14	Overlap Factor Values	79
Table 15	Radio Properties	80
Table 16	AM Modeling Radio Buttons	81
Table 17	Design Model Radio Buttons.....	82
Table 18	Floor Planning Features.....	83
Table 19	AP Property Search	93
Table 20	Sample Building.....	95
Table 21	Create a Building.....	96
Table 22	AP Configuration Function Overview	102
Table 23	AP Profiles to AP Groups	113
Table 24	Applying WLAN Profiles to AP Groups	114
Table 25	Profiles for Example Configuration.....	116
Table 26	AAA Profile Parameters	118
Table 27	Virtual AP Profile Parameters	119
Table 28	High-Throughput Radio Profile Configuration Parameters	126
Table 29	802.11k Profile Parameters	128
Table 30	RF Optimization Profile Parameters.....	129
Table 31	RF Event Profile Parameters.....	131
Table 32	20 MHz and 40 MHz Static Channel Configuration Options.....	134
Table 33	ARM Profile Types.....	141
Table 34	ARM Profile Configuration Parameters	142
Table 35	Rap Console Summary Tab Information	164
Table 36	Rap Console Connectivity Tab Information.....	166
Table 37	Remote AP Modes of Operation and Behavior	168
Table 38	Mesh Link Metric Computation.....	192
Table 39	Mesh Radio Profile Configuration Parameters.....	200
Table 40	802.11a/802.11g RF Management Configuration Parameters.....	205
Table 41	Mesh High-Throughput SSID Profile Configuration Parameters	211
Table 42	Mesh Cluster Profile Configuration Parameters	215
Table 43	RADIUS Server Configuration Parameters	232
Table 44	RADIUS Authentication Response Codes	233

Table 45	LDAP Server Configuration Parameters	234
Table 46	TACACS+ Server Configuration Parameters	235
Table 47	Windows Server Configuration Parameters.....	235
Table 48	Internal Database Configuration Parameters.....	236
Table 49	Server Rule Configuration Parameters.....	244
Table 50	Server Types and Purposes.....	246
Table 51	Authentication Timers	249
Table 52	802.1x Authentication Profile Basic WebUI Parameters	255
Table 53	Role Assignment for User and Machine Authentication	261
Table 54	VLAN Assignment for User and Machine Authentication	262
Table 55	Firewall Policy Rule Parameters	280
Table 56	User Role Parameters.....	283
Table 57	Conditions for User-Derived Role	287
Table 58	IPv4 Firewall Parameters	289
Table 59	WISPr Authentication Profile Parameters	297
Table 60	Captive Portal Authentication Profile Parameters.....	315
Table 61	Captive Portal login Pages.....	317
Table 62	VIA Connectivity Behavior	333
Table 63	VIA - Authentication Profile Parameters.....	336
Table 64	VIA - Connection Profile Options	338
Table 65	Configure VIA client WLAN profile.....	342
Table 66	Supported VPN AAA Deployments.....	348
Table 67	MAC Authentication Profile Configuration Parameters.....	369
Table 68	Control Plane Security Migration Strategies.....	374
Table 69	Control Plane Security Parameters	375
Table 70	Configure Campus AP Whitelist Parameters.....	377
Table 71	View Campus AP Whitelist Parameters	377
Table 72	View the Campus AP Whitelist via the CLI.....	378
Table 73	Control Plane Security Whitelists.....	380
Table 74	Master and Local Controller Whitelist Information	382
Table 75	CLI Commands to Display Cluster Settings	386
Table 76	Basic tasks to configure mobility domain	400
Table 77	Example entries	403
Table 78	Client Roaming Status.....	405
Table 79	User Roaming status.....	405
Table 80	IP Mobility Configuration Parameters.....	406
Table 81	Command Syntax.....	411
Table 82	VRRP Parameters	413
Table 83	VRRP Commands	416
Table 84	Port State Comparison.....	421
Table 85	Port Role Descriptions.....	422
Table 86	RSTP Default Values	423
Table 87	W-600 Series Controller by the Numbers.....	427
Table 88	Multi-function Media Eject Button	440
Table 89	IDS Profiles	466
Table 90	IDS General Profile Configuration Parameters	467
Table 91	Predefined IDS General Profiles	467
Table 92	IDS Denial of Service Profile Configuration Parameters.....	468
Table 93	Predefined IDS DoS Profiles.....	469
Table 94	IDS Rate Thresholds Profile Configuration Parameters	471

Table 95	IDS Impersonation Profile Configuration Parameters.....	472
Table 96	Predefined Signatures.....	473
Table 97	Signature Rule Attributes.....	474
Table 98	IDS Unauthorized Device Profile Configuration Parameters.....	475
Table 99	Default and Predefined IDS Unauthorized Device Profiles.....	478
Table 100	WMS Configuration Parameters.....	480
Table 101	AP Manual Classification Categories.....	482
Table 102	Valid SSIDs with Multi-Tenancy and Misconfigured AP Protection.....	483
Table 103	Management Password Policy Settings.....	500
Table 104	Allowed Characters in a Management User Password.....	501
Table 105	CSR Parameters.....	505
Table 106	Certificate Show Commands.....	506
Table 107	Imported Certificate Locations.....	507
Table 108	SNMP Parameters for the Controller.....	507
Table 109	Software Modules.....	509
Table 110	Logging Levels.....	510
Table 111	Guest Provisioning—Guest Field Descriptions.....	512
Table 112	File Transfer Configuration Parameters.....	522
Table 113	License Usage per License.....	529
Table 114	MIPS Controller AP Capacity.....	530
Table 115	IPv6 Client Authentication.....	537
Table 116	IPv6 Firewall Parameters.....	537
Table 117	IPv6 Firewall Policy Rule Parameters.....	539
Table 118	Default Voice Net Services and Ports.....	545
Table 119	Services for ALGs.....	547
Table 120	VoIP Call Admission Control Configuration Parameters.....	550
Table 121	Examples of Dial Plans.....	554
Table 122	WMM Access Category to 802.1D Priority Mapping.....	562
Table 123	WMM Access Category to DSCP Mappings.....	563
Table 124	WMM Access Categories and 802.1d Tags.....	565
Table 125	EDCA Parameters Station and EDCA Parameters AP Profile Settings.....	566
Table 126	Character-matching operators in regular expressions.....	594
Table 127	Regular expression repetition operators.....	594
Table 128	Regular expression anchors.....	595
Table 129	Configure option 60 on the Windows DHCP server.....	598
Table 130	Features not Supported in Each Forwarding Mode.....	603
Table 131	Predefined Network Services.....	604
Table 132	Predefined Policies.....	606
Table 133	Predefined Roles.....	608
Table 134	Predefined Management Roles.....	610
Table 135	Default (Trusted) Open Ports.....	613
Table 136	Web Page Authentication Variables.....	629
Table 137	Provision using Static IP.....	644
Table 138	Provision using PPPoE Connection.....	645

This User Guide describes the features supported by Dell PowerConnect ArubaOS and provides instructions and examples for configuring controllers and Access Points (APs). This chapter covers:

- “Audience” on page 37
- “Fundamentals” on page 37
- “Related Documents” on page 38
- “Conventions” on page 38
- “Contacting Support” on page 39

Audience

This guide is intended for system administrators responsible for configuring and maintaining wireless networks and assumes you are knowledgeable in Layer 2 and Layer 3 networking technologies.

Fundamentals

Throughout this document reference are made to controllers:

- MIPS Controllers— W-3000, W-600 Series Controller

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

WebUI

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

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.



Note: 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 items are part of the complete documentation for the Dell user-centric network:

- *Dell PowerConnect Controller Installation Guides*
- *Dell PowerConnect Access Point Installation Guides*
- *Dell PowerConnect ArubaOS Quick Start Guid*
- *Dell PowerConnect ArubaOS User Guide*
- *Dell PowerConnect ArubaOS Command Line Reference Guide*
- *Dell PowerConnect Release Notes*

Conventions

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

Table 1 *Typographical Conventions*

Type Style	Description
<i>Italics</i>	This style is used to emphasize important terms and to mark the titles of books.
System items	This fixed-width font depicts the following: <ul style="list-style-type: none">• 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]	In the command examples, items 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:



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



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



Warning: Indicates a risk of personal injury or death.

Contacting Support

Table 3 *Web site contact*

Web Site	
• Main Site	http://www.dell.com
• Support Site	http://www.support.dell.com
• Documentation Site	http://www.support.dell.com/manuals

This chapter describes how to connect a Dell controller and Dell APs to your wired network. After completing the tasks described in this chapter, see [“Access Points” on page 101](#) for information on configuring APs.

This chapter describes the following topics:

- [“Configuring the User-Centric Network” on page 41](#)
- [“Deployment and Configuration Tasks” on page 41](#)
- [“Configuring the Controller” on page 44](#)
- [“Configuring a VLAN for Network Connection” on page 45](#)
- [“Deploying APs” on page 49](#)
- [“Additional Configuration” on page 53](#)

Configuring the 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 [Chapter 27 on page 527](#))



Note: 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.




Note: 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.

Deployment and Configuration Tasks

This section describes typical deployment scenarios and the tasks you must perform in connecting an Dell controller and Dell APs to your wired network. For details on performing the tasks mentioned in these scenarios, refer to the remaining sections within this chapter.

Deployment Scenario #1




In this deployment scenario, the APs and controller are on the same subnetwork and will use IP addresses assigned to the subnetwork. 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 Dell Discovery Protocol (ADP) to locate the controller.

Configure the SSID(s) with VLAN 1 as the assigned VLAN for all users.

Deployment Scenario #2



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.



Note: 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



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.



Note: 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 PowerConnect 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 [Chapter 27, “Software Licenses” on page 527](#).
6. Connect the ports on the controller to your network.

This section describes the steps in detail.

Running the 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.

Note: 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.



Note: 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. Most countries impose penalties and sanctions for operators of wireless networks with devices set to improper country codes.

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 *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 [Chapter 26, “Management Access”](#) 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.



Note: 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.

Configuring a VLAN for Network Connection

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](#)” on page 43).

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.
- Assign to the VLAN the port(s) 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 and Updating a VLAN

You can create and update a single VLAN or bulk VLANs using the WebUI or the CLI. See [“Creating and Updating VLANs” on page 55](#).



Note: 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**.

Viewing Existing VLAN IDs

Use the CLI to view VLAN IDs.

```
(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

You can create, update, delete a VLAN pool using the WebUI or the CLI. See [“Creating, Updating and Deleting VLAN Pools” on page 56](#).

Adding existing VLAN IDs to a VLAN Pool in the CLI

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
-----
newgroup            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 <-- 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 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”](#) on page 43).



Note: 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.



Note: 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**. Note that 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 524](#).

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 [Chapter 27, “Software Licenses” on page 527](#) 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.



Note: In many deployment scenarios, an external firewall is situated between various Dell devices. [Appendix B, “External Firewall Configuration”](#) 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](#)” on page 42 and “[Deployment Scenario #3](#)” on page 43), ping the VLAN 1 IP address from a workstation on the network.
- If you created and configured a new VLAN (“[Deployment Scenario #3](#)” on page 43), ping the IP address of the new VLAN from a workstation on the network.

Deploying APs

Dell 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 deploy APs by doing the following steps:

1. Run the Java-based RF Plan tool to help position APs and import floorplans for your installation.
2. Ensure that the APs can locate the controller when they are connected to the network. There are several ways in which APs can locate the controller.
3. When deploying APs in a mesh networking environment, you must define the mesh cluster profile, mesh radio profile, and provision the AP as a mesh portal or mesh point. Note that this step is required only if you are configuring a mesh nodes.
4. Install the APs by connecting the AP to an Ethernet port. If power over Ethernet (PoE) is not used, connect the AP to a power source.
5. On the controller, configure the APs.

The following sections explain each of the above steps.

Running the RF Plan

The Java-based RF Plan tool is an application that allows you to determine AP placement based on your specified coverage and capacity requirements without impacting the live network. For more information about using RF Plan, see the *RF Plan Installation and User Guide*.

Enabling APs to Connect to the Controller

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 ensure the following:

- When connected to the network, each AP is assigned a valid IP address
- APs are able to locate the controller

Dell APs use Trivial File Transfer Protocol (TFTP) during the AP's 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. [Appendix B, “External Firewall Configuration”](#) describes the network ports that must be configured on the external firewall to allow proper operation of the network.

Enabling APs to Obtain IP Addresses

Each AP requires a unique IP address on a subnetwork that has connectivity to a controller. Dell recommends using the Dynamic Host Configuration Protocol (DHCP) to provide IP addresses for APs; the DHCP server can be an existing network server or a controller configured as a DHCP server.

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.

Enabling the DHCP server on the controller 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.

Enable the DHCP server on the controller in the CLI

```
(host) (config)# ip dhcp excluded-address ipaddr ipaddr2
(host) (config)# ip dhcp pool name
    default-router ipaddr
    dns-server ipaddr
    domain-name name
    network ipaddr mask
(host) (config)# service dhcp
```

Locating the Controller

An AP can discover the IP address of the controller in the following ways:

- From a DNS server
- From a DHCP server
- Using the Dell Discovery Protocol (ADP)

At boot time, the AP builds a list of controller IP addresses and then tries these addresses in order until a controller is reached successfully. The list of controller addresses is constructed as follows:

1. 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.
2. 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.
3. 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.
4. 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.

This list of controller IP addresses provides an enhanced redundancy scheme for controllers that are located in multiple data centers separated across Layer-3 networks.

From a DNS Server

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**.

For information on how to configure a host name entry on the DNS server, refer to the vendor documentation for your server.



Note: Dell recommends using 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.

When using DNS, the AP can learn 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 LMS.

From 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. APs identify themselves with a vendor class identifier set to Dell AP in their DHCP request. When the DHCP server responds to the 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 [Appendix A, “”](#) or refer to the vendor documentation for your server.

Using the Dell Discovery Protocol (ADP)

ADP is enabled by default on all Dell APs and controllers. To use ADP, all APs and controllers must be connected to the same Layer-2 network. If the devices are on different networks, a Layer-3 compatible discovery mechanism, such as DNS, DHCP, or IGMP forwarding, must be used instead.

With ADP, APs send out periodic multicast and broadcast queries to locate the master controller. You might need to perform additional network configuration, depending on whether the APs are in the same broadcast domain as the controller:

- If the APs are in the same broadcast domain as the master controller, the controller automatically responds to the APs' queries with its IP address.
- If the APs are not in the same broadcast domain as the master controller, you must 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.

To verify that ADP and IGMP join options are enabled on the controller, use the following CLI command:

```
(host) #show adp config
ADP Configuration
-----
key          value
---          -
discovery    enable
igmp-join    enable
```

If ADP or IGMP join options are not enabled, use the following CLI commands:

```
(host) (config) #adp discovery enable
(host) (config) #adp igmp-join enable
```

Provisioning APs for Mesh

The information in this section applies only if you are configuring and deploying APs in a mesh networking environment. If you are not, proceed to [“Installing APs” on page 52](#).

Before you install APs in a mesh networking environment, you must do the following:

- Define and configure the mesh cluster profile and mesh radio profile before configuring an AP to operate as a mesh node. An AP configured for mesh is also known as a mesh node.
- Provision one of the following mesh roles 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 [Chapter 7, “Secure Enterprise Mesh” on page 189](#).

Installing APs

Use the AP placement map generated by RF Plan to install APs. 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 from Dell Networks. For more information, see the *Installation Guide* for the specific AP.

Once an AP is connected to the network and powered up, it attempts to locate the master controller using one of the methods described in “[Locating the Controller](#)” on page 51.

On the master controller, you can view the APs that have connected to the controller in the WebUI. Navigate to the **Configuration > Wireless > AP Installation** window. [Figure 1](#) shows an example of this window.

Figure 1 APs Connected to Controller

AP Name	AP Group	AP IP	AP Type	AP MAC Address	AP Serial Number	Status
00:0b:86:c0:cf:d8	default	192.168.10.254	65	00:0b:86:c0:cf:d8	A90008272	Up 3h:33m:18s

Updating the RF Plan

After deploying APs, update the AP placement map in RF Plan. This allows more accurate reconciliation of location tracking features provided by the user-centric network—for example, locating users, intruders, rogue APs and other security threats, assets, and sources of RF interference—with the physical environment.

Additional Configuration

After you have installed a basic user-centric network, the APs advertise the default SSID. Wireless users can connect to this SSID but because you have not yet configured authentication, policies, or user roles, they will not have access to the network. Other chapters in this guide describe how to build upon this basic deployment to configure user roles, firewall policies, authentication, authentication servers, and other wireless features.

This chapter describes some basic network configuration on the controller. This chapter describes the following topics:

- “Configuring VLANs” on page 55
- “Configuring Ports” on page 57
- “About VLAN Assignments” on page 60
- “Configuring Static Routes” on page 65
- “Configuring the Loopback IP Address” on page 66
- “Configuring the Controller IP Address” on page 66
- “Configuring GRE Tunnels” on page 67

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 exist only 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.

Using 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 “[Create a Bulk VLANs using the WebUI](#)” on page 56 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, click the port in the **Port Selection** section.
5. Click **Apply**.

Using CLI

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

Create a Bulk VLANs using 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**.

Using CLI

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

Creating, Updating and Deleting VLAN Pools

You can create, update and delete a VLAN pool.

Creating a VLAN pool using the WebUI

The following configurations create a VLAN Pool named mygroup. VLAN IDs 2, 4 and 12 are then assigned to the VLAN pool mygroup.

1. Navigate to **Configuration > Network > VLAN**.
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. Names must be between 1 and 32 characters; spaces are not allowed. The VLAN name can not be modified; choose the name carefully.
5. 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. Or, click the drop-down list to view the IDs then click the <-- arrow to add the ID to the pool.
6. You must add two or more VLAN IDs to create a pool.
7. 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 (it has two or more IDs assigned to it), its status is enabled. If you create a VLAN pool and add only one or no VLAN IDs, its status appears as disabled.

8. Click **Apply**.
9. At the top of the window, click **Save Configuration**.

Updating a VLAN Pool

1. On the **VLAN Pool** window, click **Modify** next to the VLAN name you want to edit.
2. Modify 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**.

Create a VLAN Pool using CLI

The pool option allows you to create a VLAN pool consisting of two more VLAN IDs.

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

Viewing existing VLAN IDs using 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
1170   VLAN1170      FE1/7
```

Adding existing VLAN IDs using CLI

The following example illustrates adding 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 Name	Pool Status	VLAN IDs
newgroup	Enabled	2,4,12
group123	Disabled	

Configuring Ports

Both Fast Ethernet and Gigabit Ethernet ports can be set to access or trunk mode. By default, a port is in access mode 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. 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

By default, physical ports on the controller are trusted and are typically connected to internal networks while untrusted ports connect to third-party APs, public areas, or other networks to which access controls can be applied. 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/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 4 *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 make them untrusted, and designate a policy through which VLAN traffic on this port must pass.

Using 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. Select the policy 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**.

Using 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
```

Configure Trusted/Untrusted Ports and VLANs in Trunk Mode

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

Using 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 the ones 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. Or, to make only one VLAN untrusted, select a VLAN from the drop-down menu.
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 make *trusted*. (In this format, for example: 200-300, 401-500 and so on). Only VLANs listed in this range are trusted. Or, to make only one VLAN 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 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 VLAN and policies, click **Cancel**.
12. Click **Apply**.

Using CLI

```
(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
```

About VLAN Assignments

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

1. The default VLAN is the VLAN configured for the WLAN (see [“Virtual APs” on page 105](#)).
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 the VLAN 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. This does not require any server-derived rule.
6. After client authentication, the VLAN can be derived from Vendor Specific Attributes (VSA) for RADIUS server authentication. This does not require any server-derived rule. If a VSA is present, it overrides any previous VLAN assignment.

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 must be assigned a static IP address.

Using 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**.

Using CLI

```
interface vlan <id>
  ip address <address> <netmask>
```

Configuring a VLAN to Receive a Dynamic Address

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

- Manually configured by the network administrator. This is the default method and is described in [“Assigning a Static Address to a VLAN” on page 60](#). 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. These methods are described in the following section.

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, the controller can be connected to a DSL or cable modem, or a broadband remote access server (BRAS). [Figure 2](#) 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 on the uplink device. The DHCP server on the controller assigns IP addresses to users on the local network from a configured pool of IP addresses.

Figure 2 IP Address Assignment to VLAN via DHCP or PPPoE



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 on only one VLAN on the controller; this VLAN 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.
- Only one VLAN on the controller can obtain its IP address through DHCP or PPPoE. You cannot enable both the DHCP and PPPoE client on the controller at the same time.

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.

Using 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. Click **Apply**.

Using the CLI

```
vlan <id>
interface vlan <id>
  ip address dhcp-client
```

Enabling the PPPoE Client

To authenticate to 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.

Using 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. Click **Apply**.

Using CLI

```
ip pppoe-service-name <service-name>
ip pppoe-username <name>
ip pppoe-password <password>

vlan <vlan>
interface vlan <vlan>
    ip address pppoe
```

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.

Using the WebUI

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

Using CLI

```
ip default-gateway import
```

Configuring DNS/WINS Server from DHCP/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.

Using 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**.

Using CLI

```
ip dhcp pool employee-pool
  default-router 10.1.1.254
  dns-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.

Using 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**.

Using CLI

```
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 cause NAT to be performed 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.

Example Configuration

In the following example, the controller operates within an enterprise network. VLAN 1 is the outside VLAN. Traffic from VLAN 6 is source NATed using the IP address of the controller. In this example, 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 3 Example: Source NAT using Controller IP Address



Using 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**.

Using CLI


```
interface vlan 1
ip address 66.1.131.5 255.255.255.0
interface vlan 6
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 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 4](#), 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 4 *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 should 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 CLI

```
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:

Using 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** to add this route to the routing table. The message **Configuration Updated Successfully** confirms that the route has been added.

Using CLI

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



Note: If you are using the loopback IP address to access the WebUI, changing the loopback IP address will result in loss of connectivity. Dell recommends 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.

Using CLI

```
interface loopback ip address <address>  
write memory
```

Using the CLI to reboot the controller

Enter the following command in Enable mode:

```
reload
```

Configuring the Controller IP Address

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

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 the **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 menu. This list only contains VLAN IDs that have statically assigned IP addresses. If a loopback interface IP address has been previously configured then it will also appear in this list. Dynamically assigned IP addresses, for example 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 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 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 **aruba-master** to discover the IP address of the controller.

The controller also supports GRE tunnels between the controller and other GRE-capable devices. This section describes how to configure a GRE tunnel to such a device and how to direct traffic into the tunnel.



Note: The controller uses GRE tunnels for communications between master and local controllers; these GRE tunnels are automatically created and are not subject to the configuration described in this section.

Creating a Tunnel Interface

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

- Tunnel ID: this can be 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
- Tunnel destination: the IP address of the remote endpoint of the tunnel on the other GRE device.

Using 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 (check) **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**.

Using CLI

```
interface tunnel <id>
  tunnel mode gre ip
  ip address <ipaddr> <netmask>
  no shutdown
  tunnel source {<ipaddr>| loopback | vlan <vlan>}
  tunnel destination <ipaddr>
```

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
- 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 65](#) 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 69](#) for more information on how GRE tunnel status is determined). If you have more than one GRE tunnel configured, you can create multiple firewall policy rules with each rule redirecting the same traffic to different tunnels. If the tunnel in the first traffic redirect rule is down, then the tunnel in the subsequent traffic redirect rule is used instead.

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**.

CLI

```
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 tunnel. If you enable tunnel keepalives, the tunnel is considered to be “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.

By default, the controller sends keepalive frames at 10-second intervals and retries keepalives up to three times before the tunnel is considered to be down. You can reconfigure the intervals from the default. For the interval, specify a value between 1-86400 seconds. For the retries, specify a value between 0-1024.

Using 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 (check) **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**.

Using CLI

```
interface tunnel id
    tunnel keepalive [<interval> <retries>]
```


RF Plan is a wireless deployment modeling tool that helps you design an efficient Wireless Local Area Network (WLAN) that optimizes coverage and performance, without complicated WLAN network setup. RF Plan provides the following critical functionality:

- Defines WLAN coverage.
- Defines WLAN environment security coverage.
- Assesses equipment requirements.
- Optimizes radio resources.

RF Plan provides a view of each floor, allowing you to specify how you want to provide wireless coverage for each area. RF Plan also generates coverage maps with AP and AM placement.

Unlike other static site survey tools that require administrators to have intricate knowledge of building materials and other potential radio frequency (RF) hazards, RF Plan calibrates coverage in real-time through a sophisticated RF calibration algorithm. This real-time calibration lets you characterize the indoor propagation of RF signals to determine the best channel and transmission power settings for each AP. You can program the calibration to occur automatically or you can manually launch the calibration at any time to quickly adapt to changes in your wireless environment.

This chapter discusses the following topics:

- [“Supported Planning” on page 71](#)
- [“Before You Begin” on page 72](#)
- [“Launching the RF Plan” on page 74](#)
- [“Using the FQLN Mapper in the AP Provision Page” on page 94](#)
- [“RF Plan Example” on page 95](#)

Supported Planning

All the features included in the WebUI RF Plan tool will aid you in the planning 802.11n standard compliant deployments.

This WebUI RF Plan supports planning of the following types of deployments:

- **802.11n Deployments**—The RF Plan now supports planning of network environments that use the Dell’s AP-12x series of indoor access points, which are 802.11n compliant. RF Plan supports the planning of these APs in the following capacity: 802.11a/n, 802.11b/g/n, or 802.11a/b/g/n.
- **802.11n Hotspot Deployment within an Existing Environment**—This type of environment requires that AP/AM locations be fixed at the building level, see [“Fix All Suggested AP/AMs” on page 90](#). If you set and fix the location of APs prior to planning for the 802.11n APs, the APs will not move when you initialize/optimize the 802.11n AP locations.
- **802.11n Hotspot Deployment and New Environment**—The RF Plan allows you to plan for a new deployment that uses an 802.11n hotspot and 802.11a and/or 802.11 b/g support outside of the hotspot.

To plan for this type of deployment, start by planning your 802.11n hotspot. When you initialize and optimize the APs planned for the hotspot, the 802.11n APs are placed within the hotspot area. However, the same AP type will also be placed outside of the hotspot area with 802.11n support disabled.

RF Plan will deploy APs outside of the hotspot area based on the 802.11a and/or 802.11b/g rates defined by the system. For the system to define 802.11a and/or 802.11b/g rates, the system looks at the defined 802.11n rate and the distance covered by the defined rate; it then selects corresponding 802.11a and/or 802.11b/g rates based on the distance covered.

Before You Begin

Review the following steps to create a building model and plan the WLAN for your model.

Task Overview

1. Gather information about your building's dimensions and floor plan.
2. Determine the level of coverage you want for your APs and AMs.
3. Create a new building and add its dimensions.
4. Enter the parameters of your AP coverage.
5. Enter the parameters of your AM coverage.
6. Add floors to your building and import the floor plans.
7. Define special areas.
8. Generate suggested AP and AM tables by executing the AP/AM Plan features.

Planning Requirements

You should collect the following information before using RF Plan. Having this information readily available will expedite your planning efforts.

- Building dimensions
- Number of floors
- Distance between floors
- Number of users and number of users per AP
- Radio type(s)
- Overlap Factor
- Desired data rates for APs
- Desired monitoring rates for AMs
- Areas of your building(s) that you do not necessarily want coverage
- Areas of your building(s) where you do not want or cannot deploy an AP or AM
- Areas of your building(s) where you want to deploy an 802.11n Hotspot (Zone)
- Any area where you want to deploy a fixed AP or AM

Use the worksheets ([Table 5](#), [Table 6](#), and [Table 7](#)) to collect your information:

Table 5 *Planning Worksheet - Building Dimensions*

Building Dimensions
Height:
Width:
Number of Floors:
Number of Users:
Users per AP:
Radio Types:
AP Type:
Overlap Factor:
802.11a Desired Rate:
802.11n (HT) Support:
Use 40 MHz Channel Spacing:
802.11n Desired Rate:

Table 6 *Planning Worksheet - AP Desired Rates*

AP Desired Rates (2.4 GHz Radio Properties)
802.11b/g Desired Rate:
802.11n (HT) Support:
Use 40 MHz Channel Spacing:
802.11n Desired Rate:

Table 7 *Planning Worksheet - AM Desired Rates*

AM Desired Rates
802.11b g:
802.11a:
Don't Care/Don't Deploy Areas
802.11n Hotspot (Zone) Areas



Note: If 802.11n (HT) support is enabled, the system will automatically define the 802.11a and/or 802.11b/g rate as applicable. For details, see [“Radio Properties \(Desired Rates and HT Support Options\)” on page 80.](#)

Launching the RF Plan


This section describes how to launch the RF Plan and enter information in RF Plan windows.

To launch RF Plan from the WebUI, click the **Plan** tab in the WebUI menu bar. When you launch the RF Plan, the browser window displays the Campus List page.

Campus List Page

The Campus List is the first page you see when you start RF Plan. This list contains a default campus and any campus you have defined using the RF Plan software.

Figure 5 *Plan>Campus List Window*



You may add, edit, and delete campuses using this page. You may also import and export campus information. [Table 8](#) details the buttons on the Campus page.


Table 8 *Definition of Campus List Buttons*

Buttons	Description
New Campus	Use this button to create a new campus.
Browse Campus	Use this button to edit existing campuses from the campus list. To edit a campus, select the checkbox next to the campus name, then click Browse Campus . When you edit a campus, you can access other RF Plan pages.
Rename Campus	Use this button to rename an existing campus in the list. To rename a campus, select the checkbox next to the campus name, then click Rename Campus . A dialog box appears into which you enter the new name of the campus. Click OK to accept the new name, or click Cancel to exit this action.
Delete Campuses	Use this button to delete existing campuses in the list. To delete a campus, select the checkbox next to the building ID, then click Delete Campuses . You can only delete empty campuses. If you attempt to delete a campus that contains one or more buildings, an error message appears.
Export	Use this button to export a database file with all the specifications and background images of one or more selected campuses in the list. See “Exporting and Importing Files” on page 90.
Import	Use this button to import database files that define campuses into the RF Plan list. See “Exporting and Importing Files” on page 90.
AP FQLN Mapper	The AP name is a fully-qualified location name (FQLN) in the format <i>APname.floor.building.campus</i> (the <i>APname</i> portion of the FQLN must be unique). The FQLN is not case sensitive and supports a maximum of 249 characters, including spaces. You can use any combination of characters except a new line, carriage return, and non-printable control characters. You can manually set the FQLN for the AP by clicking the AP FQLN Mapper button. Setting the FQLN will reboot the APs. See “FQLN Mapper” on page 92

Building List Pane

Edit a campus from the building list pane.

Figure 6 *Plan>Building List Pane*



You can add, edit, and delete buildings using this page. You may also import and export building information. The buttons on this page are defined in [Table 9](#).

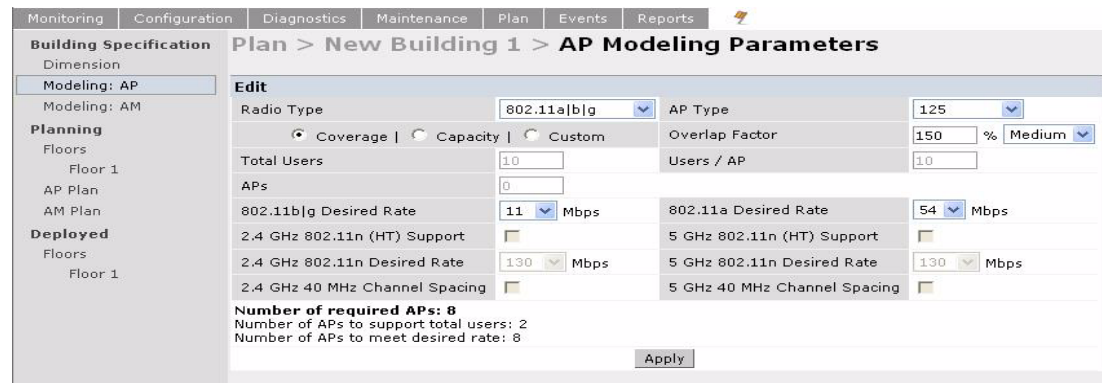
Table 9 *Building List Buttons*

Buttons	Description
New Building	Use this button to create a new building. When you add or edit a building, you can access other RF Plan pages.
Edit Building	Use this button to edit existing buildings in the building list. To edit a building, select the checkbox next to the building ID, then click Edit Building . When you add or edit a building, you can access other RF Plan pages.
Delete Buildings	Use this button to delete existing buildings in the building list. To delete a building, select the checkbox next to the building ID, then click Delete Building .
Export	Use this button to export a database file with all the specifications and background images of one or more selected buildings in the building list. See “Exporting and Importing Files” on page 90 .
Import	Use this button to import database files that define buildings into the RF Plan building list. See “Exporting and Importing Files” on page 90 .
Locate	Use this button to locate Wi-Fi devices in a building. See “Locate” on page 92 .
AP FQLN Mapper	The AP name is a fully-qualified location name (FQLN) in the format <i>APname.floor.building.campus</i> (the <i>APname</i> portion of the FQLN must be unique). The FQLN is not case sensitive and supports a maximum of 249 characters, including spaces. You can use any combination of characters except a new line, carriage return, and non-printable control characters. You can manually set the FQLN for the AP by clicking the AP FQLN Mapper button. Setting the FQLN will reboot the APs. See “FQLN Mapper” on page 92 .

Building Specifications Overview

The Building Specification Overview window displays the default values for a building that you are adding or the current values for a building that you are modifying.

Figure 7 *Plan>New Building>Overview Window*



The Overview page includes the following:

- Building Dimensions: Your building's name and dimensions
- Access Point Modeling Parameters
- Air Monitor Modeling Parameters
- **Building Dimension** button (in the upper right-hand portion of the page). Click on this button to edit the building dimensions settings.

When you create or edit information for a building, there are several ways you can navigate through RF Plan windows:

- The navigation pane on the left side of the browser window displays RF Plan pages in the order in which they should be accessed when you are creating a new building. If you are editing a building, simply click on the page you want to display or modify.
- A button for the next page appears in the upper right-hand portion of the page. You can click on this button to display the next page. For example, the **Building Dimension** button appears in the Building Specifications Overview page.
- Clicking **Apply** on editable pages sequences you to the next page. For example, when you click **Apply** in the Building Dimensions page, the AP Modeling Parameters page displays.

Building Dimension Page

The Building Dimension page allows you to specify the name and identification for the building and its dimensions. [Table 10](#) defines the parameters to insert in this window.

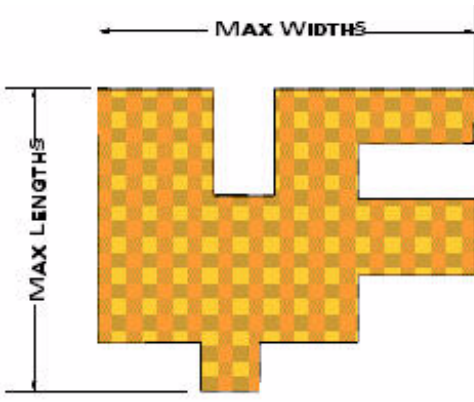
Figure 8 *Plan>New Building>Specification Window*

[Table 10](#) contains the information for you to enter in the Specification window.

Table 10 *New Building Specifications Parameters*

Parameter	Description
Campus Name	Select a campus for this building from the drop-down menu.
Building Name	The Building Name is an alphanumeric string up to 64 characters in length.

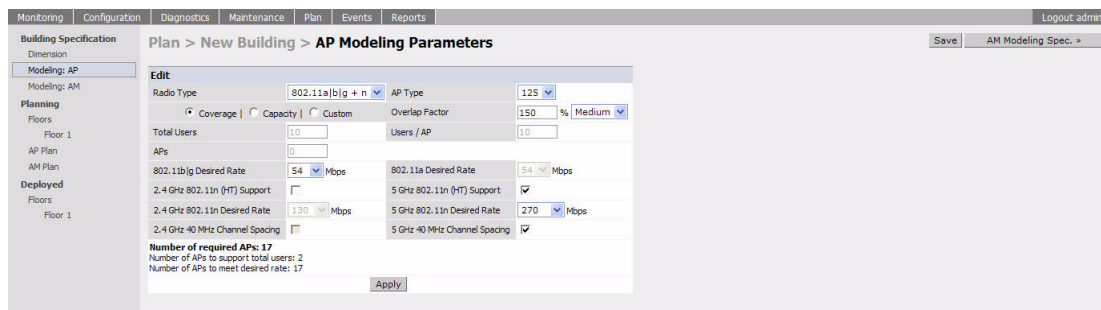
Table 10 New Building Specifications Parameters (Continued)

Parameter	Description
Width and Length	<p>Enter the rectangular exterior dimensions of the building. The valid range for this field is any integer from 1 to a value corresponding to 1x10,000.</p>  <p>If your building has an irregular shape, the width and length should represent the maximum width and length of the overall footprint of the building as seen from above. For example: When width and length are specified, RF Plan creates a rectangular area in the Planning feature pages that represent the overall area covered by the building. You need to import an appropriate background image (see “Floor Editor Dialog Box” on page 84.) to aid you in defining areas that do not require coverage or areas in which you do not wish to deploy APs and AMs (see “Area Editor Dialog Box” on page 85).</p>
Inter-Floor Height	<p>This is the distance between floor surfaces in the building. The valid range for this field is any integer from 1 to a value corresponding to 1x10,000. RF Plan uses the inter-floor height to allow APs on one floor to service users on adjacent floors. If you do not want RF Plan to factor adjacent floors, select a high inter-floor height value (for example, 300). NOTE: This is not the distance from floor to ceiling. Some buildings have a large space between the interior ceilings and the floor above.</p>
Floors	<p>Enter the number of floors in your building here. The valid range for this field is any integer from 1 to 255. A building can have a maximum of 255 floors. You can also configure negative floor IDs. Negative floor IDs let you allocate floors as sub floors, ground floors, basements or other underground floors, or floors where you do not need to deploy APs. You specify a negative integer when modifying an existing floor; you do not configure negative floor settings when adding a building or adding a floor. For more information, see “Level” on page 84.</p>
Unit	<p>Specify the unit of measurement for the dimensions you specified on the page. The choices are feet and meters.</p>

AP Modeling Parameters Page

The AP Modeling Parameters page allows you to specify the information necessary for RF Plan to determine the appropriate placement of your APs. These settings are on a per-building basis. If you have a mix of APs, choose the most common one to define the building parameters.

Figure 9 Plan>New Building>AP Modeling Parameters Window



This window allows you to select or control the parameters as defined in Table 11.

Table 11 AP Modeling Parameters

Parameter	Description
Radio Type	Use this drop-down menu to specify the radio type. See “Radio Type” on page 78
AP Type	Dell AP device. Use the drop-down menu to select the device type. The supported APs listed in the drop-down menu are dependent on the selected radio type.
Design Model	Use the Coverage, Capacity, and Custom radio buttons to specify a design model to use in the placement of APs. See “Design Model” on page 78
Overlap Factor	Use this field and drop-down to specify an overlap factor. See “Overlap Factor” on page 79 .
Users	Use this field to specify the number of users on your WLAN. See “Users/AP” on page 79 .
Radio Properties (Desired Rates and HT Support Options)	Use this drop-down to define 802.11a, 802.11b/g, and 802.11n settings for the 5 GHz and 2.4 GHz frequency bands, including high-throughput, data rates, and 40 Mhz channel spacing See “Radio Properties (Desired Rates and HT Support Options)” on page 80 .
APs	Use this field to enter the fixed number of APs to be used in this building's network (Custom model only).

Radio Type

Use the drop-down radio type menu to specify radio type of your AP. The available types are defined in [Table 12](#).

Table 12 Radio Type Definitions

Parameter	Description
801.11a/b/g	Simultaneous use of 802.11b/g and 802.11a.
802.11b/g	2.4 GHz, Direct Spread Spectrum (DSSS) multiplexing with data rates up to 11 Mbps, combined with Orthogonal Frequency Division Multiplexing/Complementary Code Keying (OFDM/CCK) with data rates up to 54 Mbps.
802.11a	5 GHz Orthogonal Frequency Division Multiplexing (OFDM) with data rates up to 54 Mbps.
802.11a/b/g + n	Mixed-mode radio type which allows for simultaneous use of 802.11b/g and 802.11n traffic on the 2.4 GHz frequency band, and 802.11a and 802.11n traffic on the 5 GHz frequency band.
802.11b/g + n	Mixed-mode radio type that allows for simultaneous use of 802.11b/g and 802.11n traffic on the 2.4 GHz frequency band.
802.11a + n	Mixed-mode radio type that allows for simultaneous use of 802.11a and 802.11n traffic on the 5 GHz frequency band.

Design Model

Three radio buttons, defined in [Table 13](#), allow you to control the kind of model used to determine the number and type of APs.

Table 13 Design Model Radio Buttons

Radio Button	Description
Coverage	Use this option to let RF Plan automatically determine the number of APs based on desired data rates and the configuration of your building. The higher the data rate, the smaller the coverage area, and the more APs that are required. Coverage is the most common type of installation.
Capacity	Use this option to let RF Plan determine the number of APs based on the total number of users, ratio of users to APs, and desired data rates. Capacity-based coverage is useful for high capacity conference or training rooms, where the APs could have a high volume of users.
Custom	Use this option to specify a fixed number of APs. Custom coverage is useful for deployments with a known number of APs or if you have a fixed project budget.

Overlap Factor

The Overlap Factor is the amount of signal area overlap when the APs are operating. Overlap is important if an AP fails as it allows the network to self-heal with adjacent APs powering up to assume some of the load from the failed device. Although there may be no holes in coverage in this scenario, there is likely to be a loss of throughput. Increasing the overlap allows for higher throughputs when an AP has failed and allows for future capacity as the number of users increases.

You can select a pre-determined value from the drop-down overlap menu or specify a value in the text box to the left of the drop-down. The following table describes the available options.

Table 14 Overlap Factor Values

Overlap Factor	Description
100% Low	Use this option for buildings that contain open spaces such as warehouses.
150% Medium	Use this option for most typical office environments with cubicles and sheetrock walls that have higher WLAN user density than warehouses.
200% High	Use this option for dense deployments such as buildings with poor RF coverage characteristics including buildings with thick brick or concrete walls, lots of metal, or excess RF noise (for example, data centers).
Custom	Use this option to enter a custom rate. For most office spaces, 120% works well. When specifying the custom rate, the valid range is 1% to 1000%.

Users/AP



Note: The Users text boxes are active only when the Capacity model is selected.

Enter the number of users you expect to have on your WLAN in the Users text box. Enter the number of users per AP you expect in the Users/AP text box.

The numbers entered in the these two text boxes must be non-zero integers between 1-255 inclusive.

Radio Properties (Desired Rates and HT Support Options)

Define 802.11a, 802.11b/g, and 802.11n settings for the 5 GHz and 2.4 GHz frequency bands, including high-throughput, data rates, and 40 Mhz channel spacing.

Table 15 *Radio Properties*

Radio Property	Description
802.11a Desired Rate	<p>The desired 802.11a rate defines the estimated transmit rate within the WLAN coverage area. The higher the speed, the smaller the coverage area, and the more APs required. The valid values are: 54, 48, 36, 24, 18, 12, 9, 6.</p> <p>This option is only available when 802.11n (HT) support is disabled (unchecked or grayed out). When an 802.11n radio type, such as 802.11a + n or 802.11a/b/g + n, is selected and 802.11n (HT) support is enabled (checked) on the 5 GHz band, the system will automatically define the 802.11a rate. The system looks at the defined 802.11n rate and the distance covered by the defined rate; the system then selects a corresponding 802.11a rate based on the distance covered.</p>
5 GHz 802.11 (HT) Support	<p>High-throughput is available when utilizing the IEEE 802.11n standard and can be enabled on the 5 GHz frequency band when either the 802.11a + n or 802.11a/b/g + n mixed-mode radio type is selected.</p> <p>The 802.11n (high-throughput) draft standard supports MIMO (Multiple Input, Multiple Output) and the option of 40 MHz mode of operation. However, high-throughput can be utilized on a 20 MHz channel or on a 40 MHz channel (bonded channel pair).</p>
5 GHz 802.11n Desired Rate	<p>The desired 802.11n rate defines the estimated transmit rate within the WLAN coverage area. The higher the speed, the smaller the coverage area, and the more APs required.</p> <p>This option is only available when 802.11n (HT) support is enabled (checked).</p> <p>The valid values when using 20 MHz channel spacing: 6.5, 13.0, 19.5, 26.0, 39.0, 52.0, 58.5, 65.0, 78.0, 104.0, 117.0, 130.0.</p> <p>The valid values when using 40 MHz channel spacing: 13.5, 27.0, 40.5, 54.0, 81.0, 108.0, 121.15, 135.0, 162.0, 216.0, 243.0, 270.0.</p>
5 GHz Use 40 MHz Channel Spacing	<p>Use 40 MHz Channel Spacing—40 MHz operation, which supports higher data rates by utilizing two 20 MHz channels as a bonded pair, requires that high-throughput be enabled (checked). 40 MHz mode is most often utilized on the 5 GHz frequency band due to a greater number of available channels.</p> <p>This option is only available when 802.11n (HT) support is enabled (checked).</p>
802.11b/g Desired Rate	<p>The desired 802.11b/g rate defines the estimated transmit rate within the WLAN coverage area. The higher the speed, the smaller the coverage area, and the more APs required. The valid values are: 54, 48, 36, 24, 18, 12, 9, 6, 11, 5.5, 2, 1.</p> <p>This option is only available when 802.11n (HT) support is disabled (unchecked or grayed out). When an 802.11n radio type, such as 802.11g + n or 802.11a/b/g + n, is selected and 802.11n (HT) support is enabled (checked) on the 2.4 GHz band, the system will automatically define the 802.11b/g rate. The system looks at the defined 802.11n rate and the distance covered by the defined rate; the system then selects a corresponding 802.11b/g rate based on the distance covered.</p>
2.4 GHz 802.11 (HT) Support	<p>High-throughput is available when utilizing the IEEE 802.11n standard and can be enabled on the 2.4 GHz frequency band when either the 802.11g + n or 802.11a/b/g + n mixed-mode radio type is selected.</p> <p>The 802.11n (high-throughput) draft standard supports MIMO (Multiple Input, Multiple Output) and the option of 40 MHz mode of operation. However, high-throughput can be utilized on a 20 MHz channel or on a 40 MHz channel (bonded channel pair).</p>
2.4 GHz 802.11n Desired Rate	<p>The desired 802.11n rate defines the estimated transmit rate within the WLAN coverage area. The higher the speed, the smaller the coverage area, and the more APs required.</p> <p>This option is only available when 802.11n (HT) support is enabled (checked).</p> <p>The valid values when using 20 MHz channel spacing: 6.5, 13.0, 19.5, 26.0, 39.0, 52.0, 58.5, 65.0, 78.0, 104.0, 117.0, 130.0.</p> <p>The valid values when using 40 MHz channel spacing: 13.5, 27.0, 40.5, 54.0, 81.0, 108.0, 121.15, 135.0, 162.0, 216.0, 243.0, 270.0.</p>

Table 15 Radio Properties (Continued)

Radio Property	Description
2.4 GHz Use 40 MHz Channel Spacing	40 MHz operation, which supports higher data rates by utilizing two 20 MHz channels as a bonded pair, requires that high-throughput be enabled (checked). Due to a limited number of channels on the 2.4 GHz frequency band, 40 MHz mode is most often utilized on the 5 GHz frequency band where a greater number of channels are available. This option is only available when 802.11n (HT) support is enabled (checked).


AM Modeling Page

The AM Modeling page allows you to specify the information necessary for RF Plan to determine the appropriate placement of your AMs.



Note: AM coverage rates refer to the rate at which an AM captures packets. RF Plan uses that information to determine the placement of AMs.

Figure 10 AM Modeling Page



Controls on this page allow you to select the following functions, which are described in more detail in this section:

Table 16 AM Modeling Radio Buttons

Radio Button	Description
Design Model	Use these radio buttons to specify a design model to use in the placement of AMs. See “Design Models” on page 82 .
Monitor Rates	Use this drop-down menu to specify the desired monitor rate for the AMs. See “Monitor Rates” on page 82 .
AMs	Use this field to manually specify the number of AMs to deploy (Custom Model only).

Design Models

Two radio buttons on the page allow you to specify the model used to determine the number and type of APs.

Table 17 Design Model Radio Buttons

Radio Button	Description
Coverage	Use this option to let RF Plan automatically determine the number of AMs based on desired monitor rates and the configuration of the building. Desired rate is selectable from 1 to 54 Mbps in the Coverage model.
Custom	Use this option to specify a fixed number of AMs. When the AM Plan portion of RF Plan is executed, RF Plan distributes the AMs evenly.

Note: The monitor rates you select for the AMs should be less than the data rates you selected for the APs. If you set the rate for the AMs at a value equal to that specified for the corresponding PHY type AP, RF Plan allocates one AM per AP. If you specify a monitor rate greater than the data rate, RF Plan allocates more than one AM per AP.

Monitor Rates


Use the drop down menus to select the desired monitor rates for the 2.4 GHz (802.11b/g) and 5 GHz (802.11a) frequency bands. The available monitor rates that display in drop-down lists will vary: these rates are dependent on the radio type selected on AP modeling page and they will also be adjusted to accommodate for 20 MHz vs. 40 MHz channel spacing when 802.11n (HT) support is enabled.

Note: This option is available only when the coverage design model is selected.

Planning Floors Page



The Planning Floors page enables you to see the footprint of your floors.

Figure 11 Planning Floors Page



You can select or adjust the features as described in [Table 18](#).

Table 18 Floor Planning Features

Feature	Description
Zoom	Use this drop-down menu or type a zoom factor in the text field to increase or decrease the size of the displayed floor area. See “Zoom” on page 83 .
Approximate Coverage Map (select radio type)	Use this drop-down to select a particular radio type for which to show estimated coverage. See “Approximate Coverage Map” on page 83 .
Edit Floor	Click on this link to launch the Floor Editor dialog box. See “Floor Editor Dialog Box” on page 84 .
New in Areas section	Click on this link to launch the Area Editor dialog box. See “Area Editor Dialog Box” on page 85 .
New in Suggested Access Points and Air Monitors section	Click on this link to launch the Suggested Access Point Editor dialog box. See “Access Point Editor Dialog Box” on page 86 .
Status in Deployed Access Points and Air Monitors section	<p>The Status column displays the status of each AP for the floor you are viewing within a live network.</p> <p>Up: AP is up (live). The corresponding AP icon on the floor map will display a live AP icon.</p>  <p>Down: AP is down. The corresponding AP icon on the floor map will display with a red “X” over the AP icon symbolizing that the AP is down.</p> 

Zoom

The Zoom control sets the viewing size of the floor image. It is adjustable in finite views from 10% to 1000%. You may select a value from the drop-down zoom menu or specify a value in the text box to the left of the drop-down. When you specify a value, RF Plan adjusts the values in the drop-down to display a set of values both above and below the value you typed in the text box.

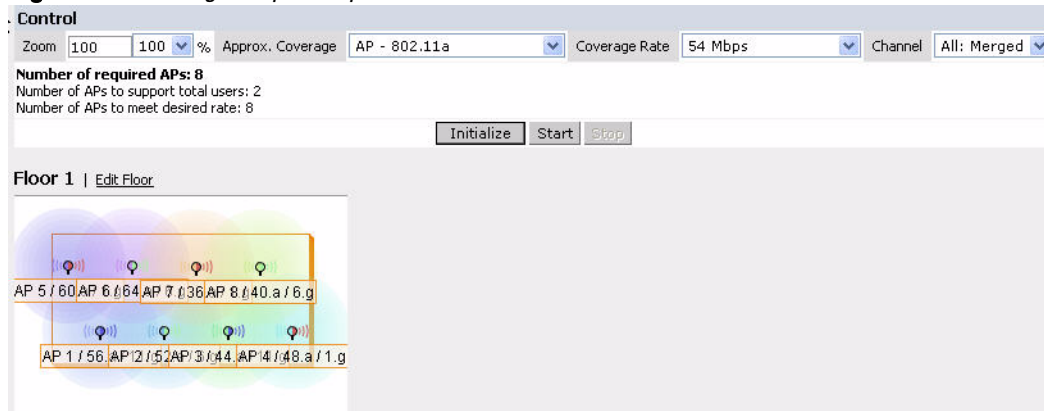
Approximate Coverage Map

Select a radio type from the **Coverage** drop-down menu to view the approximate coverage area for each of the APs that RF Plan has deployed in AP Plan or AM Plan. Adjusting the coverage values help you to understand how the AP coverage works in your building.



Note: You will not see coverage areas displayed here until you have executed either an AP Plan or an AM Plan.


Figure 12 Coverage Map Example



Floor Editor Dialog Box

The Floor Editor dialog box allows you to modify the floor level, specify the background image, and name the floor. The Floor Editor is accessible from the Floors Page by clicking on the **Edit Floor** link.

Figure 13 Floor Editor Dialog Box



Level

When modifying an existing floor, you can configure it with a negative integer to specify a basement or some other underground floor that you do not need or want to deploy APs.

To configure a negative floor, specify a negative integer in the Level field. The valid range is -100 to 255; however, a building can have a maximum of 255 floors.

Naming

You may name the floor anything you choose as long as the name is an alphanumeric string with a maximum length of 64 characters. The name you specify appears to the right of the Floor Number displayed above the background image in the Planning view.

Background Images

You can import a background image (floor plan image) into RF Plan for each floor. A background image is extremely helpful when specifying areas where coverage is not desired or areas where an AP/AM is not to be physically deployed.

Use the guidelines in this section when importing background images. By becoming familiar with these guidelines, you can ensure that your graphic file is edited properly for pre- and post-deployment planning.

- **Edit the image**—Use an appropriate graphics editor to edit the file as needed.
- **Scale the image**—If the image is not scaled, proportional triangulation and heat map displays can be incorrect when the plan is deployed.
- **Calculate image dimensions**—Calculate the image pixels per feet (or meters) against a known dimension. Use that value to calculate the width and length of the image.
- **Leave a border around the image**—When creating the image, leave a boarder around the image to help triangulate Wi-Fi devices outside of the building.

- **Multiple floors**—If your building has multiple floors, make sure there is a common anchor point for all floors; for example an elevator shaft, a staircase, and so on.
- **Larger dimensions**—Use larger dimensions only for scaling to more accurately calculate the full dimensions. For best results, final floor images 2048 X 2048 and smaller perform best.

Select a background image using the Browse button on the Floor Editor dialog box.

- **File Type and Size**

Background images must be JPEG format and may not exceed 2048 X 2048 pixels in size. Attempting to import a file with a larger pixel footprint than that specified here results in the image not scaling to fit the image area in the floor display area.

Because background images for your floors are embedded in the XML file that defines your building, you should strongly consider minimizing the file size of the JPEGs that you use for your backgrounds. You can minimize the file size by selecting the maximum compression (lowest quality) in most graphics programs.

Note: The WebUI displays floor plans using Adobe Flash Player, which does not support progressive JPEG images. If you have a progressive JPEG image you want to use as background image, open the image in an image editing program and re-save the image with standard/baseline compression.

- **Image Scaling**

Images are scaled (stretched) to fit the display area. The display area aspect ratio is determined by the building dimensions specified on the Dimension page.


Area Editor Dialog Box

The Area Editor dialog box allows you to specify areas on your building floors where you either do not care about coverage, or where you do not want to place an AP or AM.

Open the Area Editor dialog box by clicking **New** in the Areas section.

You specify these areas by placing them on top of the background image using the Area Editor.

Figure 14 Area Editor Dialog Box



Naming

Logical name of area, as an alphanumeric string consisting of 1 to 64 characters. Dell recommends that you provide a meaningful name to the area to ensure that it is readily identifiable.

Location and Dimensions

Specify absolute coordinates for the lower left corner and upper right corner of the box that represents the area being defined.

- Begin the measurement with the lower left corner of the rectangular display area that represents your building's footprint.
- The coordinates of the upper right-hand corner of the display area are the absolute values of the dimensions you provided for the building.

Location settings are zero-based. Values range from 0 to (height -1 and width -1). For example, coordinates of the upper right corner for a building that measures 200 ft. wide x 400 ft. in length, would be 199 and 399.



Note: The unit of measurement displayed as either feet or meters is based on your building settings. See [“Building Dimension Page” on page 76](#) for details about configuring building parameters.

You may also use the drag and drop feature of the Area Editor to drag your area to where you want it and resize it by dragging one or more of the handles displayed in the corners of the area.

Area Types



Select one of the area types from the drop-down menu: **Don't Care**, **Don't Deploy**, or **802.11n Zone**.

- **Don't Care:** Coverage is not required in the area specified in this dialog box. This specification typically applies to areas where coverage cannot be guaranteed.
This setting results in the display of an orange rectangle at the associated area in the floor diagram.
- **Don't Deploy:** No APs are to be positioned in the area specified in this dialog box.
This setting results in display of a yellow rectangle at the associated area in the floor diagram.
- **802.11n Zone:** 802.11n compliant APs are required to be positioned in the area specified in this dialog box only. This setting results in display of a green rectangle at the associated area in the floor diagram.

You cannot right-click within an existing area to add another area inside of it. For instance, if a Don't Care or Don't Deploy Area needs to overlap with an 802.11n Zone, you must create each of the areas outside of one another and then move them to the correct position of overlap. You can click and drag the areas to the appropriate positions of overlap, or you can right-click on the area to modify its location.

Access Point Editor Dialog Box

The Access Point Editor allows you to manually create or modify a suggested AP.

To create an AP, open the Access Point Editor dialog box by clicking **New** in the Suggested Access Points and Air Monitors section.

To modify an existing AP, place the cursor over the AP and click it to display the Suggested Access Point Editor dialog box.

Figure 15 Access Point Editor

Suggested Access Point Editor			
Name	AP 11	Floor Name	Floor 1
Fixed	No	Radio	802.11a b g
X	100	Y	50
802.11b g Type	Access Point		
802.11b g Channel	1	802.11b g Power Level	14.0 dBm
2.4 GHz 802.11n (HT) Support	<input type="checkbox"/>	Use 2.4 GHz 40 MHz Channel	<input type="checkbox"/>
802.11a Type	Access Point		
802.11a Channel	36	802.11a Power Level	14.0 dBm
5 GHz 802.11n (HT) Support	<input type="checkbox"/>	Use 5 GHz 40 MHz Channel Pair	<input type="checkbox"/>
Memo	<div style="border: 1px solid #ccc; height: 40px;"></div>		
<input type="button" value="Apply"/> <input type="button" value="Cancel"/> <input type="button" value="Delete"/>			

Naming

RF Plan automatically names APs using the default convention `ap number`, where *number* starts at 1 and increments by one for each new AP. When you manually create an AP, the new AP is assigned the next number and is added to the bottom of the suggested AP list.

You may name an AP anything you wish. The name must consist of alphanumeric characters and be 64 characters or less in length.

Fixed

Fixed APs do not move when RF Plan executes the positioning algorithm.



Note: You might typically set a fixed AP when you have a specific room, such as a conference room, in which you want saturated coverage. You might also want to consider using a fixed AP when you have an area that has an unusually high user density.

Choose **Yes** or **No** from the drop-down menu. Choosing **Yes** locks the position of the AP as it is shown in the coordinate boxes of the Access Editor. Choosing **No** allows RF Plan to move the AP as necessary to achieve best performance.

Radio Types

The Radio drop-down menu allows you to specify what frequency band the AP uses. You can choose from one of the following:


- 802.11a/b/g (2.4 GHz and 5 GHz frequency bands)
- 802.11a (5 GHz frequency band)
- 802.11 b/g (2.4 GHz frequency band)



Note: 802.11n (HT) support features are available on the 2.4 or 5 GHz frequency band. The availability of these options on these frequency bands is dependent on the radio (frequency band) chosen and whether or not these feature were enabled on the AP modeling page at the building level.

X and Y Coordinates

The physical location of the AP is specified by X-Y coordinates that begin at the lower left corner of the display area. The numbers you specify in the X and Y text boxes are whole units. The Y-coordinate increases as a point moves up the display and the X-coordinate increases as they move from left to right across the display.



802.11 Types

The 802.11 b/g and 802.11a Type drop-down menus allow you to choose the mode of operation for the AP. You may choose to set the mode of operation to Access Point or Air Monitor.

802.11 Channels

The 802.11a and 802.11b/g channel drop-down menus allow you to select from the available channels.



Note: The available channels vary depending on the regulatory domain (country) in which the device is being operated.

802.11 Power Levels

The power level drop-down menus allow you to specify the transmission power of the AP. Choices are OFF, 0, 1, 2, 3, and 4. A setting of 4 applies the maximum Effective Isotropic Radiated Power (EIRP) allowed in the regulatory domain (country) in which you are operating the AP.

802.11n Features

- **802.11n (HT) Support (2.4 or 5 GHz):** Specify if 802.11n high-throughput support should be enabled on this AP.

In order to enable high-throughput on a new AP being added to the plan at the floor level, 802.11n (HT) support must first be enabled at the building level within the AP modeling parameters. If not, this option will be grayed out. See [“AP Modeling Parameters Page” on page 77](#) for details about AP modeling parameters.

- **Use 40 MHz Channel (2.4 or 5 GHz):** Specify if 802.11n high-throughput support should utilize a 40 MHz channel (bonded channel pair).

In order to select a valid 40 MHz channel for a new AP being added at the floor level, use of 40 MHz channel spacing must first be enabled at the building level within the AP modeling parameters. If not, this option will be grayed out. See [“AP Modeling Parameters Page” on page 77](#) for details about AP modeling parameters.

If high-throughput is enabled and use of a 40 MHz channel pair is not enabled, a 20 Mhz channel will be utilized.


Memo

The **Memo** text field allows you to enter notes regarding the AP. You can enter a maximum of 256 alphanumeric characters in the **Memo** field.

AP Plan Page

The AP Plan page uses the information entered in the modeling pages to locate APs in the building(s) you described. All of the options on the Floors page can also be viewed and configured on the AP Plan page. The AP Plan page also includes some additional options, such as initializing, optimizing, and fixing AP/AM locations.

Figure 16 AP Planning



Initialize

Initialize the Algorithm by clicking the **Initialize** button. This makes an initial placement of the APs and prepares RF Plan for the task of determining the optimum location for each of the APs. As soon as you click **Initialize** you see the AP symbols appear on the floor plan.

Colored circles around the AP symbols on the floor plan indicate the approximate coverage of the individual AP and the color of the circle represents the channel on which the AP is operating. The circles appear when you select an *approximate coverage* value on one of the Floors pages. You may also click an AP icon and drag it to manually reposition it.

Optimize

Click **Optimize** to launch the optimizing algorithm. The AP symbols move on the page as RF Plan finds the optimum location for each.

The process may take several minutes. You may watch the progress on the status bar of your browser. The algorithm stops when the movement is less than a threshold value calculated based on the number of APs. The threshold value may be seen in the status bar at the bottom of the browser window.

Viewing the Results

The results of optimizing algorithm may be viewed two ways: graphically and in a table of suggested APs. You may obtain information about a specific AP by placing the cursor over its symbol. An information box appears that contains information regarding location, radio type, high-throughput support, channel(s), and power.

The Suggested Access Points and Air Monitors table lists the coordinates, power, location, power setting, high-throughput support, and channel(s) for each of the APs that are shown in the floor plan.

Fix All Suggested AP/AMs

Fix existing AP/AM locations at the building level. If AP/AM locations are fixed, AP/AMs will not move from their fixed locations during initialization or optimization. Clicking on this button will fix the locations of existing APs and AMs. You only need to click this button on either the AP or AM Plan page.

AM Plan Page

The AM Plan page uses the information entered in the modeling pages to locate AMs in the building(s) you described and calculate the optimum placement for the AMs. All of the options on the Floors page can also be viewed and configured on the AM Plan page. The AM Plan page also includes some additional options, such as initializing, optimizing, and fixing AP/AM locations.

Initialize

Initialize the Algorithm by clicking **Initialize**. This makes an initial placement of the AMs and prepares RF Plan for the task of determining the optimum location for each of the AMs. When you click **Initialize**, the AM symbols appear on the floor plan.

Optimize

Click **Optimize** to launch the optimizing algorithm. The AM symbols move on the page as RF Plan finds the optimum location for each.

The process may take several minutes. You may watch the progress on the status bar of your browser. The algorithm stops when the movement is less than a threshold value calculated based on the number of AMs. The threshold value may be seen in the status bar at the bottom of the browser window.

Viewing the Results

Viewing the results of the AM Plan feature is similar to that for the AP Plan feature.

The results of optimizing algorithm may be viewed two ways: graphically and in a table of suggested AMs. You may obtain information about a specific AM by placing the cursor over its symbol. An information box appears that contains information about the exact location, PHY type, high-throughput-support, channel, power, and so on.

The Suggested Access Points and Air Monitors table lists the coordinates, power, location, power setting, and channel for each of the AMs that are shown in the floor plan.

Fix All Suggested AP/AMs

Fix existing AP/AM locations at the building level. If AP/AM locations are fixed, AP/AMs will not move from their fixed locations during initialization or optimization. Clicking on this button will fix the locations of existing APs and AMs. You only need to click this button on either the AP or AM Plan page.

Exporting and Importing Files

Both the Campus List page and the Building List page have **Export** and **Import** buttons, which allow you to export and import files that define the parameters of your campus and buildings. You can export a file so that it may be imported into and used to automatically configure a controller. On a controller, you can import a file that has been exported from another controller or from the standalone version of RF Plan that runs as a Windows application.



Note: The WebUI version of RF Plan only supports JPEG file formats for background images.

The files that you export and import are XML files and, depending on how many buildings are in your campus, floors are in your buildings, and how many background images you have for your floors, the XML files may be quite large. (See [“Background Images” on page 84.](#))



Note: In order for the WebUI RF Plan tool to import and read a standalone plan that incorporates 802.11n standard APs and was originally created in the Java-based standalone RF Plan tool, the plan must be exported out from the standalone tool using the Controller WebUI Format (version 3.0).

Export Campus

To export a file that defines the parameters of one or more campuses, including all of its associated buildings, select the campus(es) to be exported in the Campus List page and then click **Export**.

After you click the **Export** button, you are prompted to include the background images.



When exporting a campus file, Dell recommends that you click **OK** to export the background images. If you click **Cancel**, the exported file does not include the background images. The File Download window appears.

From the File Download window, click **Save** to save the file. The Save As dialog box appears. From here, navigate to the location where you want to save the file and enter the name for the exported file. When naming your exported file, be sure to give the file the .XML file extension, for example, *My_Campus.XML*.

Exported campus files include detailed information about the campus and the selected building(s).

Import Campus

You can import only XML files exported from another controller or from the standalone version of RF Plan that runs as a Windows application.




Note: Importing any other file, including XML files from other applications, may result in unpredictable results.

To import a file that defines the building parameters of one or more campuses, click the **Import** button in the Campus List page. The Import Buildings page appears, as described in [“Import Buildings Page” on page 92](#).

Export Buildings Page

To export a file that defines the parameters of one or more buildings, select the building(s) to be exported in the Building List page and then click **Export**.

After you click the Export button, you are prompted to include the background images.



When exporting a building file, Dell recommends that you click **OK** to export the background images. If you click **Cancel**, the exported file does not include the background images. The **File Download** window appears.

From the File Download window, click **Save** to save the file. The Save As dialog box appears. From here, navigate to the location where you want to save the file and enter the name for the exported file. When naming your exported file, be sure to give the file the .XML file extension, for example, *My_Building.XML*.

Exported building files include the name of the campus to which the building belongs; however, detailed campus parameters are not included.


Import Buildings Page

You can import only XML files exported from another controller or from the standalone version of RF Plan that runs as a Windows application.



Note: Importing any other file, including XML files from other applications, may result in unpredictable results.

To import a file that defines the parameters of one or more buildings, click the **Import** button in the Building List page.



In the Import Buildings page, click **Browse** to select the file to be imported, then click the **Import** button.

Locate

The **Locate** button on the Building List page allows you to search for APs, AMs, monitored clients, etc. on a building by building basis. To use this feature, select the building in which you want to search, and click **Locate**.

The Target Devices table displays information on each of these devices. To add a device, click **Add Device**. To delete a device, click **Remove Device**. To select a device, click **Choose Devices**.

FQLN Mapper

Both the Campus List page and the Building List page have the **AP FQLN Mapper** button, which allows you to create a fully-qualified location name (FQLN) for the specified AP/AM in the format *APname.Floor.Building.Campus*.

The FQLN is not case sensitive and supports a maximum of 249 characters, including spaces. You can use any combination of characters except a new line, carriage return, and non-printable control characters.

You can use the FQLN mapper for multiple purposes, including:

- Searching for deployed APs/AMs
- Configuring the AP name in the form APname.Floor.Building.Campus
- Modifying the location of APs

To use this feature, select one or more campuses from the Campus List page, or one or more buildings from the Building List page, and click **AP FQLN Mapper**.

The AP FQLN Mapper page appears. From here, you can search for deployed APs by entering one or more parameters in the Search fields, view the results in the Search Results table, configure the FQLN, and modify the location of an AP.

To search for deployed APs, enter information in the Search fields and click **Search**.

The screenshot shows a web interface titled "Plan > AP FQLN Mapper". Below the title is a "Search" section with a "Hide Search" link. The search form includes the following fields:

- AP Name: [text input]
- FQLN: [text input]
- Wired MAC: [text input]
- Serial Number: [text input]
- IP Address: [text input]
- Status: [dropdown menu with 'Any' selected]

 At the bottom of the search section, there is a "Number of results per page:" dropdown set to "10", and two buttons: "Search" and "Reset".

You can perform a search based on one or more of the following AP properties:

Table 19 AP Property Search

Property	Description
AP Name	Logical name of the AP or AM. You can enter a portion of the name to widen the search.
Wired MAC	MAC address of the AP or AM. You can enter a portion of the MAC address to widen the search.
IP Address	IP address of the AP or AM. You can enter a portion of the IP address to widen the search.
FQLN	Fully-qualified location name of the AP, in the form APname.floor.building.campus. You can enter a portion of the FQLN to widen the search.
Serial Number	Serial number of the AP. You can enter a portion of the serial number to widen the search.
Status	Current state of the AP, including Up/Down/Any.

Use the drop-down list to the right of the Number of results per page to specify the number of APs to display in the search results. After entering the search criteria, you can either click **Reset** to clear the entries or click **Search** to search for APs. If you click **Search**, the results are displayed in the Search Result table.

You can view the information in ascending or descending order. By default, the display is in ascending order, based on the AP name (the white arrow indicates the row that is being used to sort the information). Left-click on a column head to view the information in ascending or descending order (you may need to click multiple times to get the desired display).

In addition to displaying AP names, wired MAC addresses, serial numbers, IP addresses, FQLNs, and AP status, the Search Result table displays the AP type and when it was last updating. From here you can modify the attributes that create the FQLN for the selected AP, using the following drop-down lists:

- **Campus**—Displays the campus where the AP is deployed. To deploy the AP in a different campus, select a campus from the drop-down list. The Campus defines the buildings and floors displayed.



Note: This drop-down list only displays the existing campuses that you are managing. To add a new campus, see [“Campus List Page” on page 74](#).

- **Building**—Displays the building where the AP is deployed. To deploy the AP in a different building, select a building from the drop-down list.



Note: This drop-down list only displays the available buildings in the selected campus. To add a new building, see [“Building List Pane” on page 75](#).

- **Floor**—Displays the floor where the AP is deployed. To deploy the AP on a different floor, select a floor from the drop-down list.



Note: This drop-down lists only displays the available floors in the selected building. To add a new floor, see [“Planning Floors Page” on page 82](#).

To submit your changes, click **Set FQLN**. Setting the FQLN reboots the APs.

Using the FQLN Mapper in the AP Provision Page

The AP Provision page (available from **Configuration > Wireless > AP Installation**) allows you to set an FQLN during the AP provisioning process.

Scroll to the FQLN Mapper near the bottom of the AP Provision page to modify the following attributes that create the FQLN:

- Campus
- Building
- Floor

The AP name appears in the AP List at the bottom of the page and will be used when provisioning the AP. To rename an AP, enter the new name in the AP Name field.

To retain the old FQLN value when reprovisioning an AP, *do not* select the **Overwrite FQLN** checkbox. However, if you configure new values for the campus, building, and floor settings, the FQLN value is changed, even if the **Overwrite FQLN** checkbox is selected. To remove a previously configured value, you can select N/A for a specific attribute.

If you provision more than one AP, the selected value for the campus, building, and floor is based on the first selected AP and applies to all APs. Only the AP name will be different as each AP must have a unique name.

Using the WebUI

1. Navigate to the **Configuration > Wireless > AP Installation** page. The list of discovered APs appears in the page.
2. Select the AP you want to set an FQLN, and click **Provision**.

3. Modify the FQLN attributes:
 - In the Provisioning page, scroll to the FQLN Mapper near the bottom of the page and modify the campus, building, and floor attributes.
 - Optionally, if you want rename an AP, scroll to the AP List at the bottom of the page and enter the new name in the AP Name field. For more information about AP names, see [Chapter 4, “Access Points”](#)
4. Click **Apply and Reboot**.

Using CLI

Reprovisioning the AP causes it to automatically reboot. When configuring the FQLN, you may also provision other AP settings.

The following example assumes you are not renaming an AP. For more information about AP names, see [Chapter 4, “Access Points”](#).

```
provision-ap
  read-bootinfo ap-name <name>
  copy-provisioning-params ap-name <name>
  fqln <name>
  reprovision ap-name <name>
```

RF Plan Example

This section guides you through the process of creating a building and populating it with APs and AMs using RF Plan. Ensure you have sample.JPEG floor images handy for walking through this planning example.

Sample Building

[Table 20](#) list the information to be used in this coverage-based planning example.

Table 20 *Sample Building*

Building Dimensions
Height: 100
Width: 100
Number of Floors: 2
User Information
Number of Users: N/A
Users per AP: N/A
Radio Types: 802.11a/b/g/n
AP Type: 93
Overlap Factor: 150% (Medium)
AP Desired Rates (5 GHz Radio Properties)
802.11a Desired Rate: 48 Mbps
802.11n (HT) Support: N/A
Use 40 MHz Channel Spacing: N/A
802.11n Desired Rate: N/A

Table 20 *Sample Building (Continued)*

Building Dimensions	
AP Desired Rates (2.4 GHz Radio Properties)	
802.11b/g Desired Rate:	48 Mbps
802.11n (HT) Support:	N/A
Use 40 MHz Channel Spacing:	N/A
802.11n Desired Rate:	N/A
AM Desired Rates	
802.11b g:	24 Mbps
802.11a:	24 Mbps
Don't Care/Don't Deploy Areas	
Shipping & Receiving = Don't Care	Lobby = Don't Deploy
802.11n Hotspot (Zone) Areas	
N/A	

Create a Building

In this section you create a building using the information supplied in the planning table.

1. In the Campus List, select **New Campus**. Enter: **My Campus** and click **OK**.
2. In the Campus List, select the checkbox next to **My Campus**, and click **Browse Campus**.
3. Click **New Building**. The Overview page appears.
4. Click **Save**. A dialog box appears that indicates the new building was saved successfully. Click **OK** to close the dialog box.
5. Click **Building Dimension**. The Specification page appears.
6. Enter the following information in the text boxes.

Table 21 *Create a Building*

Text Box	Information
Campus Name	My Campus (The name is automatically populated based on what you entered in step 1)
Building Name	My Building
Width	100
Length	100
Inter Floor Height	20
Units	Feet
Floors	2

7. Click **Save**. A dialog box appears that asks if you want to save and reload this building now since the building name was changed. Click **OK** to accept.
Another dialog box appears stating that the building was saved successfully. Click **OK** to close the dialog box.
8. Click **Apply**. RF Plan returns you to the Overview page.

Model the Access Points

You now determine how many APs are required to cover your building with a specified data transfer rate and overlap.

In this example, you use the Coverage Model. The following are assumed about the performance of the WLAN:

- Radio Types: 802.11a/b/g
 - AP Type: 93
 - Overlap factor: Medium (150%)
 - 802.11a desired rate: 48 Mbps
 - 802.11b desired rate: 48 Mbps
1. From the navigation tree, Click on **Modeling:AP** under Building Specification. The AP Modeling Parameters page appears.
 2. Select 801.11 a|b|g|n from the **Radio Type** drop-down menu.
 3. Select **Medium** from the **Overlap Factor** drop-down menu.
 4. Notice that the percentage show at the left of the drop-down menu changes to 150%.
 5. Select **48** from the 802.11 b|g **Desired Rate** drop-down menu.
 6. Select **48** from the 801.11 a **Desired Rate** drop-down menu.
 7. Click **Save**, then **OK**.
 8. Click **Apply**. RF Plan moves to the AM Modeling Parameters page.

Model the Air Monitors

You now determine how many AMs are required to provide a specified monitoring rate. In this example you continue to use the Coverage Model and make the following assumptions:

- 802.11 b|g monitor rate: 24 Mbps
 - 802.11 a monitor rate: 24 Mbps
1. Select **24** from the 802.11 b|g Monitor Rate drop-down menu.
 2. Select **24** from the 802.11 a Monitor Rate drop-down menu.
 3. Click **Save**, then **OK**.
 4. Click **Apply**. RF Plan moves to the Planning page.

Add and Edit a Floor

You now add floor plans to your floors. In this section you:

- Add a background image floor plan for each floor
- Name the floors



Note: The information in this section assumes that you have a JPEG file that you can use as a sample background image when re-creating the steps.

Adding the background image and naming the first floor

1. In the Planning page, click the **Edit Floor** link at the right of the Floor 1 indicator. The Floor Editor dialog box appears.
2. Enter: **Entrance Level** in the **Name** box of the Floor Editor Dialog.
3. Use the **Browse** button to locate the background image for the 1st floor.
4. Click **Apply**.

Adding the background image and naming the second floor

1. Click the **Edit Floor** link at the right of the Floor 2 indicator.
2. Enter: **Second Level** in the **Name** box of the Floor Editor Dialog.
3. Use the **Browse** button to locate the background image for the 2nd floor.
4. Click **Apply**.
5. Click **Save** on the Planning page, then **OK**.

Defining Areas

Before you advance to the AP and AM Planning pages, define special areas, such as Don't Care, Don't Deploy, or 802.11n Zone. This example includes a Don't Care and a Don't Deploy Area.

This example assumes the following:

- We do not care if we have coverage in the Shipping and Receiving Area
- We do not want to deploy APs or AMs in the Lobby Area

Creating a Don't Care Area



Note: You can zoom in on the floor plan using the Zoom drop-down near the top of the AP Planning page, or type a zoom value in the text box at the left of the drop-down and press the enter key on your keyboard. For example, enter a zoom factor of 400.

1. In the Planning page, click the **New** link in the Areas section under Floor 1 (named Entrance Level).

This opens the Area Editor.

2. Enter: **Shipping and Receiving** in the **Name** text box in the Area Editor.
3. Select **Don't Care** from the **Type** drop-down menu box.
4. Click **Apply**.

Notice that an orange box appears near the center of the floor plan.

5. Use your mouse (or other pointing device) to place the cursor over the box.

Notice that the information you typed in the editor appears in the box. You see the name and type of area, as well as the coordinates of the lower left corner and upper right corner of the box.



Note: The $x = 0$ and $y = 0$ coordinates correspond to the lower left corner of the layout space.

6. Using your mouse, left-click and drag the box to the area of your floor plan that will represent the shipping and receiving area.
7. To position the Don't Care box, drag one corner of the box to a corresponding corner and using one of the corner handles of the box, stretch it to fit.
You can also position the box by entering values in the **Left**, **Bottom**, **Right**, and **Top** fields.
8. Click **Save**, then **OK**.

Creating a Don't Deploy Area

1. Click the **New** link in the **Areas** section under **Floor 1** (named **Entrance Level**) to open the **Area Editor**.
2. Enter: **Lobby** in the **Name** text box in the **Area Editor**.
3. Select **Don't Deploy** from the **Type** drop-down menu box.
4. Click **Apply**.
Notice that a yellow box appears near the center of the floor plan.
5. Use your mouse (or other pointing device) to place the cursor over the box.
Notice that the information you typed in the editor appears in the box. You see the name and type of area, as well as the coordinates of the lower left corner and upper right corner of the box.



Note: The $x = 0$ and $y = 0$ coordinates correspond to the lower left corner of the layout space.

6. Using your mouse, left-click and drag the box to the area of your floor plan that you wish to designate as the **Lobby Area**.
7. To position the **Don't Deploy** box, drag one corner of the box to a corresponding corner and using one of the corner handles of the box, stretch it to fit.
You can also position the box by entering values in the **Left**, **Bottom**, **Right**, and **Top** fields.
8. Click **Save**, then **OK**.

Running the AP Plan

In this section you run the algorithm that searches for the best place to put the APs.

1. From the navigation tree, click **AP Plan** under the **Planning** section. The **AP Planning** page appears.
You might want to zoom in on the floor plan. Zoom in using the zoom drop-down near the top of the **AP Planning** page, or type a zoom factor in the text box at the left of the drop-down and press the enter key on your keyboard.
Try entering a zoom factor of 400.
Notice that the number of required APs displays towards the top of the page, which represents the same value that you saw when you modeled your APs on the **AP Modeling Parameters** page. Notice that the APs are not yet displayed on the floor plan.
2. Click **Initialize**.
You should see the required total number of AP symbols appear on the two floor diagrams. Also notice that the **Suggested Access Points** tables below each floor diagram have been populated with information about the suggested APs for each corresponding floor.
3. Click **Optimize**.
After you **Initialize** the APs you must optimize the algorithm. The APs move around on the floor plans as the algorithm is running.
The algorithm stops when the movement is less than a threshold value calculated based on the number of APs. The threshold value may be seen in the status bar at the bottom of the browser window.



Note: To see the approximate coverage areas of each of the APs, select an AP type from the **Approx. Coverage** drop-down box and select a rate from the **Coverage Rate** drop-down box.

4. Click **Save**, then **OK**.

Running the AM Plan

Running the AM Plan algorithm is similar to running the AP Plan.

1. From the navigation tree, click **AM Plan** under the Planning section. The AM Planning page appears.
2. Click **Initialize** then **Optimize**.

The algorithm stops when the movement is less than a threshold value calculated based on the number of AMs. The threshold value may be seen in the status bar at the bottom of the browser window.

3. Click **Save**, then **OK**.

Dell APs receive their configuration from their host controller. At power on, an AP locates its host controller and the AP's configuration is "pushed" from the controller to the AP. This chapter describes how to configure your controllers so that your APs performs the functions required for your network.



Note: In a network with a master and local controllers, an AP will initially connect to the master controller. Alternatively, you can instruct your AP to download its configuration (and ArubaOS) from a local controller (see [Chapter 18, "Adding Local Controllers"](#) for details).

Included in the chapter is:

- "Remote AP vs Campus AP" on page 101
- "Basic Configuration" on page 102
- "AP Names and Groups" on page 102
- "Virtual APs" on page 105
- "Configuring Profiles" on page 105
- "Profile Hierarchy" on page 110
- "Virtual AP Configurations" on page 115
- "Configuring High-throughput on Virtual APs" on page 125
- "Advanced Configuration Options" on page 127
- "Automatic Channel and Transmit Power Selection Using ARM" on page 135
- "APs Over Low-Speed Links" on page 135
- "AP Redundancy" on page 136
- "AP Maintenance Mode" on page 137
- "Managing AP LEDs" on page 138

Remote AP vs Campus AP

When to use Remote AP (RAP) versus a Campus AP (CAP).

- When the network between the AP and controller is an un-trusted/non-routable network, such as the Internet, a RAP is recommended; in cases where the AP needs to connect over private links (LAN, WAN, MPLS), a CAP is recommended. The reason that CAP is not recommended over a non-routable network is because the IPsec within control plane security is in tunnel mode.
- RAP supports internal DHCP server; CAP does not.
- For both RAP or CAP, tunneled SSIDs will be brought down eight (8) seconds after the AP detects that there is no connectivity to the controller. For CAP bridge-mode SSIDs, the CAP will be brought down after the keepalive times out (default 3.5 minutes). RAP bridge mode SSIDs are configurable to stay up indefinitely (always-on / persistent).
- ARM operates on both RAPs and CAPs.
- Backup mode SSID is supported on the RAP only.

Basic Configuration

You configure APs using the WebUI and the CLI. [Table 22](#) list the basic configuration functions and features.

Table 22 AP Configuration Function Overview

Features and Function	Description
Wireless LANs	<p>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.</p> <p>The WebUI includes a WLAN Wizard that provides easy-to-follow steps to configure a new WLAN.</p> <p>NOTE: All new WLANs are associated with the ap-group named “default”.</p>
AP operation	<p>An Dell AP can function as an air monitor (AM) performing network and radio frequency (RF) monitoring. 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.</p>
Quality of Service (QoS)	<p>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.</p>
RF management	<p>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.</p> <p>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.</p>
Intrusion Detection System (IDS)	<p>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.</p>
Mesh	<p>Configure Dell APs as mesh nodes to bridge multiple Ethernet LANs or extend wireless coverage. A mesh node is either</p> <ul style="list-style-type: none">• a mesh portal—an AP that uses its wired interface to reach the controller• or a mesh point—an AP that establishes a path to the controller via the mesh portal <p>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 802.11radio to carry mesh traffic as well as WLAN services. Chapter 7, “Secure Enterprise Mesh” on page 189 contains more specific information on the Mesh feature.</p>



Note: The 802.11n features, such as high-throughput and 40 MHz configuration settings, are supported on APs that are 802.11n standard compliant; Dell’s W-AP120 Series, W-AP90 Series.

AP Names and Groups

In the user-centric network, each AP has a unique name and belongs to an AP group.

AP Names

Each AP is identified with an automatically-derived name.

- the name is the AP’s Ethernet MAC address in the format `xx:xx:xx:xx:xx:xx`

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”.



Note: Renaming an AP requires a reboot of the AP before the new name takes effect. Therefore, if you need to do this, there should be little or no client traffic passing through the AP.

In RF Plan or RF Live, the AP name can be part of a fully-qualified location name (FQLN) in the format *APname.floor.building.campus*. The *APname* portion of the FQLN must be unique.

Duplicate Names

You can display the status of APs in your database by executing the **show ap database long** command. The output will flag an AP that has a duplicate name (N flag).

To clear the AP with the duplicate name (assuming it is no longer connected to your network), use the command **clear gap-db wired-mac**.

Renaming in the WebUI

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**.

Renaming in 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.

AP Groups

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 17](#)).

Figure 17 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 procedure describes how to create an AP group and, because all discovered APs initially belong to the AP group named "default", how to reassign an AP to your newly-created AP group.



Note: Reassigning an AP from an AP group requires a reboot of the AP for the new group assignment to take effect. Therefore, if you need to do this, there should be little or no client traffic passing through the AP.

Creating an 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 name appears in the Profile list.

Assigning APs to an AP group 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.

Creating an AP group 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 an AP to an AP group 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>
```



Note: 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.


Virtual APs

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.

You can configure virtual APs 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 (see [Figure 18](#)).

Figure 18 Virtual AP Configurations Applied to the same AP



Configuring Profiles

In ArubaOS, related configuration parameters are grouped into a *profile* that you can apply as needed to an AP group or to individual APs. This section lists each category of AP profiles that you can configure and apply to an AP group or to an individual AP, and describes how the profiles are interrelated.

Each of the profiles described can be configured via the CLI or the WebUI. To see a full list of profiles available in ArubaOS, select the **Configuration** tab in the WebUI and navigate to **Advanced Services>All Profiles**. The **All Profiles** arranges group configuration profiles into six categories:

- [“Wireless LAN Profiles” on page 106](#)
- [“AP Profiles” on page 107](#)
- [“QoS Profiles” on page 108](#)
- [“RF Management Profiles” on page 108](#)
- [“IDS Profiles” on page 109](#)
- [“Mesh Profiles” on page 109](#)
- [“Controller Profiles” on page 110](#)

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 [“Configuring 802.11k Protocol” on page 127](#).

- **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.



Caution: Configuring the 802.11a and 802.11g beacon rates should only be used in conjunction 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 none of the APs in your Mesh deployment are 802.11n-capable, 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 affect immediately; rebooting is not required.
- **Virtual AP profile**—This profile defines your WLAN by enabling or disabling the bandsteering, 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.



Note: 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.

- **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 [“AAA Profile Parameters” on page 118](#).

- **XML API server profile**—Specifies the IP address of an external XML API server.
- **RFC 3576 server**—Specifies the IP address of a RFC 3576 RADIUS server.
- **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.
- **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 [Chapter 12, “Captive Portal” on page 299](#).
- **802.1x authentication profile**—Defines default user roles for machine or 802.1x authentication, and parameters for 802.1x termination and failed authentication attempts. For a list of the basic parameters in the 802.1x authentication profile, refer to [Chapter 9, “802.1x Authentication” on page 251](#)
- **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 232](#)
- **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 234](#).
- **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 235](#).
- **Server group**—This 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 244](#)
- **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 66 on page 348](#).
- **Management authentication profile**—Enables or disables management authentication, and identifies the default role for authenticated management clients. This profile also references a server group.
- **Wired authentication profile**—This profile merely references an AAA profile to be used for wired authentication.
- **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.
- **Stateful NTLM authentication Profile**—Monitor the NTLM (NT LAN Manager) authentication messages between clients and an authentication server. If the client successfully authenticates via an NTLM authentication server, the switch can recognize that the client has been authenticated and assign that client a specified user role

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.
- **Regulatory domain**—Defines the AP’s country code and valid channels for both legacy and high-throughput 802.11a and 802.11b/g radios.
- **Wired AP profile**—Controls 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). This profile also configures the switching mode characteristics for the port, and sets the port as either trusted or untrusted.
- **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.
- **Wired Port Profile**—Allows you to specify a AAA profile for users connected to the wired port on an AP.
- **AP Provisioning profile**—Defines a group of provisioning parameters for an AP or AP group.
- **AP Authorization Profile**—Allows you to assign an to a provisioned but unauthorized AP to a AP group with a restricted configuration profile.
- **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 566.
- **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 566.

QoS Profiles

The **QoS** profiles configure traffic management and VoIP functions.

- **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 “[The VoIP Call Admission Control Profile](#)” on page 550.
- **Traffic management profile**—Specifies the minimum percentage of available bandwidth to be allocated to a specific SSID when there is congestion on the wireless network, and sets the interval between bandwidth usage reports.

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.
- **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.

- **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 [Chapter 5, “Adaptive Radio Management \(ARM\)”](#) on page 139.
- **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.)
- **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.

IDS Profiles

These profiles configure the AP’s Intrusion Detection System (IDS) features, which detect and disable rogue APs and other devices that can potentially disrupt network operations. 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.

The following IDS profiles and their parameters are described in detail in “IDS Configuration” on page 466.

- **IDS General profile**—Configures AP attributes.
- **IDS Rate Thresholds profile**—Defines thresholds assigned to the different frame types for rate anomaly checking.
- **IDS signature matching profile**—Configures signatures for intrusion detection. This profile can include predefined signatures or signatures that you configure.
- **IDS DoS profile**—Configures traffic anomaly settings for Denial of Service attacks.
- **IDS impersonation profile**—Configures anomaly settings for impersonation attacks.
- **IDS unauthorized device profile**—Configures detection for unauthorized devices. Also configures rogue AP detection and containment.
- **IDS profile**—This profile manages a complete set of IDS profile parameters by referencing all other types of IDS profiles

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.
- **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.
- **Mesh cluster profile**—Contains the mesh cluster name (MSSID), authentication methods, security credentials, and cluster priority.

Controller Profiles

These controller profiles set the management password policy, define equipment OUIs, or configure VIA authentication and connection settings.

- **Valid Equipment OUI Profile**—Set one or more Dell OUIs for the controller.
- **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.
- **Dialplan Profile**—Define SIP dial plans on the controller to provide outgoing PSTN calls.

Profile Hierarchy

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:

- define a session policy
- define your server group
- create an AAA profile that references the session policy and your server group

[Figure 19](#) represents the AP and AP Group profile hierarchy in the WebUI (navigate to **Configuration>AP configuration**).

Figure 19 AP Specific and AP Group Profile Hierarchies





Figure 20 displays how the Layer 2 authentication profiles and Layer 3 authentication profiles reference other types of profiles. To view the profile hierarchy for Layer 2 authentication profiles in the WebUI, navigate to the **Configuration>Authentication** and select the **L2 Authentication** tab. To view the profile hierarchy for Layer 3 authentication profiles, navigate to **Configuration>Authentication** and select the **L3 Authentication** tab.

Figure 20 Layer 2/Layer3 Profile Hierarchies



The example below follows the suggested order of steps to configure a virtual AP.

```

vlan 60
!
ip access-list session THR-POLICY-NAME-WPA2
  user any any permit
!
user-role THR-ROLE-NAME-WPA2
  session-acl THR-POLICY-NAME-WPA2
!
  
```

```

aaa authentication dot1x "THR-DOT1X-AUTH-PROFILE-WPA2"
    termination enable
!
aaa server-group "THR-DOT1X-SERVER-GROUP-WPA2"
    auth-server Internal
!
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"
!
wlan ssid-profile "THR-SSID-PROFILE-WPA2"
    essid "THR-WPA2"
    opmode wpa2-aes
!
wlan virtual-ap "THR-VIRTUAL-AP-PROFILE-WPA2"
    ssid-profile "THR-SSID-PROFILE-WPA2"
    aaa-profile "THR-AAA-PROFILE-WPA2"
    vlan 60
!
ap system-profile "THR-AP-SYSTEM-PROFILE"
    lms-ip 1.1.1.1
    bkup-lms-ip 2.2.2.2
!
ap-group "THRQ1-STANDARD"
    virtual-ap "THR-VIRTUAL-AP-PROFILE-WPA2"
    ap-system-profile "THR-AP-SYSTEM-PROFILE"

```

Applying Profiles

You can use the “default” named profile or create a new profile that you can edit as required. You can also change the values of any parameter in a profile. ArubaOS gives you 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.

For example, if your wireless network includes a master controller in Edmonton, Alberta and a local controller in Toronto, Ontario, you can segregate the APs into two AP groups: “default” for the APs in Edmonton and “Toronto” for the APs in Toronto. The primary difference between the APs in Edmonton and Toronto is the controller from which the APs boot. The APs in Edmonton boot from the master controller, while the APs in Toronto boot from the local controller.

You configure the address of the local controller in the AP system profile. You need two instances of the AP system profile: one for Edmonton and one for Toronto, in our example. You can apply the “default” profiles for other AP profile types to both AP groups (see [Table 23](#)).

Table 23 AP Profiles to AP Groups

AP Profiles	“default” AP Group	“Toronto” AP Group
802.11a	“default”	“default”
802.11b/g	“default”	“default”
Wired	“default”	“default”
Ethernet 0 Link	“default”	“default”
Ethernet 1 Link	“default”	“default”
AP System	“default”	“Toronto”

Table 23 AP Profiles to AP Groups

AP Profiles	"default" AP Group	"Toronto" AP Group
Regulatory Domain	"default"	"default"
SNMP	"default"	"default"



Note: Each instance of a profile must have a unique name. In the example above, there are two different AP system profiles, therefore each instance should have a unique name.

You can apply the same virtual AP profiles to the AP groups shown in [Table 23](#). For example, there are users in both Edmonton and Toronto that access the same "Corpnet" WLAN. Note that if your WLAN requires authentication to an external server, you may want to have users who associate with the APs in Toronto authenticate with their local servers. In this case, you can configure a slightly different AAA profiles; one that references authentication servers in the Edmonton and the other that references servers in Toronto ([Table 24](#)).


Table 24 Applying WLAN Profiles to AP Groups

WLAN Profiles	"default" AP Group	"Toronto" AP Group
Virtual AP	"Corpnet-E"	"Corpnet-T"
SSID	"Corpnet"	"Corpnet"
AAA	"E-Servers"	"T-Servers"

When 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 virtual AP profile. You can apply multiple virtual AP profiles to individual APs, as well as to AP groups.

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.

Figure 21 Excluding a Virtual AP Profile from an AP



Excluding a virtual AP profile from an 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**.


Excluding a virtual AP profile from an AP in the CLI

```
ap-name <name>  
  exclude-virtual-ap <profile>
```

Viewing Profile Errors

You can view the list of profile errors using the WebUI or the CLI. If you have profiles with errors, the WebUI displays them with a *flag* icon next to main horizontal menu (Figure 22). Click on the flag to view the list of profile errors. A pop-up is displayed with the list of errors. Click on the error to view the configuration screen with the profile error.

Figure 22 Profile Errors



Use the `show profile-errors` command to view a list of profile errors.

```
#show profile-errors  
Invalid Profiles  
-----  
Profile                               Error  
-----  
ap regulatory-domain-profile "default" Country RU does not match controller country US  
wlan virtual-ap "test-vap"           VLAN 1000 does not exist
```

Virtual AP Configurations

This section includes examples of how to configure virtual APs for the “default” AP group, which includes all APs discovered by the controller, and for a specific AP. The example configuration contain the following WLANs:

- An 802.11a/b/g SSID called “Corpnet” that uses WPA2 and is available on all APs in the network
- An 802.11a/b/g SSID called “Guest” that uses open system and is only available on the AP “building3-lobby” (this AP will support both the “Corpnet” and “Guest” SSIDs)

Each WLAN requires a different SSID profile that maps into a separate virtual AP profile. For the SSID “Corpnet”, which will use WPA2, you need to configure an AAA profile that includes 802.1x authentication and an 802.1x authentication server group.

Because all APs discovered by the controller belong to the AP group called “default”, you assign the virtual AP profile that contains the SSID profile “Corpnet” to the “default” AP group. For the “Guest” SSID, you configure a new virtual AP profile that you assign to the AP named “building3-lobby”. [Table 25](#) list the profiles that you need to modify or create for these examples.

Table 25 Profiles for Example Configuration

AP Group/Name	Virtual AP Profile	SSID Profile	AAA Profile
“default”	“corpnet” <ul style="list-style-type: none"> • VLAN: 1 • SSID profile: “corpnet” • AAA profile: “corpnet” 	“corpnet” <ul style="list-style-type: none"> • SSID: Corpnet • WPA2 	“corpnet” <ul style="list-style-type: none"> • 802.1x authentication default role: “employee” • 802.1x authentication server group: “corpnet” <ul style="list-style-type: none"> - Radius1 - Radius2
“building3-lobby”	“guest” <ul style="list-style-type: none"> • VLAN: 2 • Deny Time Range • SSID profile: “guest” • AAA profile: “default-open” 	“guest” <ul style="list-style-type: none"> • SSID: Guest • Open system 	“default-open” (This is a predefined, read-only AAA profile that specifies open system authentication)

Configuring the WLAN

In this example Corpnet WLAN, users are validated against a corporate database on a RADIUS authentication server before they are allowed access to the network. Once validated, users are placed into a specified VLAN (VLAN 1 in this example) and assigned the user role “employee” that permits access to the corporate network.



Note: Dell recommends that you assign a unique name to each virtual AP, SSID, and AAA profile that you modify. In this example, you use the name “corpnet” to identify each of the profiles.

Follow the step below to configure the Corpnet WLAN.

1. Configure a policy for the user role **employee** and configure the user role **employee** with the specified policy.
2. Configure RADIUS authentication servers and assign them to the **corpnet** 802.1x authentication server group.
3. Configure authentication for the WLAN.
 - a. Create the **corpnet** 802.1x authentication profile.
 - b. Create the AAA profile **corpnet** and specify the previously-configured **employee** user role for the 802.1x authentication default role.
 - c. Specify the previously-configured **corpnet** 802.1x authentication server group.
4. For the AP group “default”, create and configure the virtual AP **corpnet**.
 - a. Create a new virtual AP profile **corpnet**.
 - b. Select the previously-configured **corpnet** AAA profile for this virtual AP.
 - c. Create a new SSID profile **corpnet** to configure “Corpnet” for the SSID name and WPA2 for the authentication.

Configuring the User Role

In this example, the **employee** user role allows unrestricted access to network resources and is granted only to users who have been successfully authenticated with an external RADIUS server. You can configure a more restrictive user role by specifying allowed or disallowed source and destination, protocol, and service for the traffic. For more information about configuring user roles, see [“User Roles” on page 283](#).

In the WebUI

1. Navigate to the **Configuration > Security > Access Control > Policies** page.
2. Click **Add** to add a new policy. Enter the name of the policy.

Default settings for a policy rule permit all traffic from any source to any destination, but you can make a rule more restrictive. You can also configure multiple rules; the first rule in a policy that matches the traffic is applied. Click **Add** to add a rule. When you are done adding rules, click **Apply**.

3. Click the **User Roles** tab. Click **Add** to add a new user role. Enter the name of the role. Under Firewall Policies, click **Add**. In the Choose from Configured Policies drop-down list, select the policy you previously created. Click **Done**.
4. Click **Apply**.

In the CLI

```
ip access-list session <policy>
    <source> <dest> <service> <action>
user-role employee
    access-list session <policy>
```

Configuring Authentication Servers

This example uses RADIUS servers for the client authentication. You need to specify the hostname and IP address for each RADIUS server and the shared secret used to authenticate communication between the server and the controller. After configuring authentication servers, assign them to the **corpnet** server group, an ordered list of the servers to be used for 802.1x authentication.

For more information about configuring authentication servers, see [“Configuring Servers” on page 232](#).

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > Servers** page.
2. Select **Radius Server** to display the Radius Server List.
3. Enter the name of the server, and click **Add**. The server name appears in the list of servers.
4. Select the server name. Enter the IP address and shared secret for the server. Select the **Mode** checkbox to activate the authentication server.
5. Click **Apply** to apply the configuration.
6. Select **Server Group** on the Servers page.
7. Enter the name of the group, and click **Add**. The server group name appears in the list of server groups.
8. Select the server group name. Click **New** to add a server to the group. Under Server Name, select the server you just configured and click **Add**.
9. Click **Apply** to apply the configuration.

In the CLI

```
aaa authentication-server radius Radius1
    host <ipaddr>
    key <key>
    enable
aaa server-group corpnet
    auth-server Radius1
```

Configuring Authentication

In this example, you create the 802.1x authentication profile **corpnet**. 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.

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > L2 Authentication** page. Select 802.1x Authentication Profile.
 - a. In the 802.1x Authentication Profile list on the right window pane, enter **corpnet** in the entry blank at the bottom of the list, and click **Add**.
 - b. Select the **corpnet** 802.1x authentication profile you just created.
 - c. You can configure parameters in the **Basic** or **Advanced** tabs. These parameters are described in detail in [Table 52 on page 255](#). For this example, you use the default values, so click **Apply**.
2. Select the **AAA Profiles** tab.
 - a. Scroll down to the bottom of the AAA Profiles Summary pane, then click **Add**. An entry blank appears.
 - b. Enter **corpnet**, then click **Add**.
 - c. Scroll back up the AAA Profiles Summary pane, and select the **corpnet** AAA profile you just created.
 - d. For this example, change the 802.1x Authentication Default Role, select the **employee** role you previously configured. You can also configure other the AAA profile parameters (see [Table 26](#)).
 - e. Click **Apply**.

Table 26 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 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.

3. Select the 802.1x Authentication Profile under the **corpnet** AAA profile to reveal the 802.1X Authentication Profile pane.
 - a. Click the **802.1X Authentication Profile** drop-down list and select **corpnet**.
 - b. Click **Apply**.

4. Select the 802.1x Authentication Server Group under the corpnet AAA profile to reveal the 802.1X Authentication Server Group pane.
 - a. Click the **802.1X Authentication Server Group** drop-down list and select the **corpnet** server group you previously configured.
 - b. Click **Apply**.

In the CLI

```
aaa authentication dot1x corpnet
aaa profile corpnet
  authentication-dot1x corpnet
  dot1x-default-role employee
  dot1x-server-group corpnet
```

Applying the Virtual AP

In this example, you apply the **corpnet** virtual AP to the “default” AP group which consists of all APs.

In the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration > AP Group** page.
2. Click **Edit** for the “default” AP group.
3. Select **Wireless LAN** (under **Profiles**), then select **Virtual AP**.
4. Select **New** from the **Add a profile** drop-down menu. Enter the name for the virtual AP profile (for example, **corpnet**), 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 “Dell-ap” ESSID. You must configure a new ESSID and SSID profile for the virtual AP profile before you apply the profile.

In the Profile Details entry for the new virtual AP profile, navigate to the **AAA Profile** drop-down list and select the AAA profile you previously configured to reveal the AAA Profile pop-up window. Click **Apply** to set the AAA profile and close the pop-up window.

Creating a new SSID Profile

1. Select **New** from the **SSID Profile** drop-down menu in the Profile Details entry for the new virtual AP profile. This launches an SSID profile pop-up window.
2. Enter the name for the SSID profile (for example, **anynet**).
3. Enter a name in the Network Name (SSID) field (for example, **Corpnet**).
4. Select **WPA2** for Network Authentication,.
5. Click **Apply** to set the SSID profile and close the pop-up window.
6. Click **Apply** again at the bottom of the Profile Details window.
7. Click the new Virtual AP name in the Profiles list or the Profile Details to display the configuration parameters defined in [Table 27](#).
8. Verify that **Virtual AP enable** is selected; select 1 for the VLAN.
9. Click **Apply**.

Table 27 *Virtual AP Profile Parameters*

Parameter	Description
Virtual AP enable	Select the Virtual AP enable checkbox to enable or disable the virtual AP.

Table 27 Virtual AP Profile Parameters

Parameter	Description
Allowed band	<p>The band(s) on which to use the virtual AP:</p> <ul style="list-style-type: none"> • 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.
VLAN	<p>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, then click the arrow button to associate that VLAN with the virtual AP profile.</p>
Forward mode	<p>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.</p> <p>Click the drop-down list to select one of the following forward modes:</p> <ul style="list-style-type: none"> • 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 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 de-capsulates 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. High-throughput APs in decrypt-tunnel mode do not support de-aggregation of MAC Service Data Units (A-MSDUs). <p>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.</p>
Deny time range	<p>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 virtual AP profile.</p>
Mobile IP	<p>Enables or disables IP mobility for this virtual AP. Default: Enabled</p>
HA Discovery on-association	<p>If enabled, all clients of a virtual AP will receive mobility service on association. Default: Disabled</p>
DoS Prevention	<p>If enabled, APs ignore deauthentication frames from clients. This prevents a successful deauth attack from being carried out against the AP. This does not affect third-party APs. Default: Disabled</p>

Table 27 Virtual AP Profile Parameters

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 deauth attacks. Default: Enabled
Blacklist Time	Number of seconds that a client is quarantined from the network after being blacklisted. Default: 3600 seconds (1 hour)
Multicast Optimization for Video	Enable/Disable dynamic multicast optimization. This parameter is disabled by default, and cannot be enabled without the PEFNG license.
Multicast Optimization Threshold	Maximum number of high-throughput stations in a multicast group beyond which dynamic multicast optimization stops. Range: 2-255 stations Default: 6 stations.
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.
Multi Association	Enables or disables multi-association for this virtual AP. When enabled, this feature allows a station to be associated to multiple APs. If this feature is disabled, when a station moves to new AP it will be de-authorized by the AP to which it was previously connected, deleting station context and flushing key caching information. Important things to know when using the Multi Association feature: <ul style="list-style-type: none"> • When enabled, the system allows multiple associations per client. If the maximum number of clients allowed per AP is limited to a small number there is a risk of increased association failures. • If a client has multiple associations, it may not do active scanning before roaming event which could result in it not being associated to nearest AP. • Multiple associations may result in more frequent roaming.
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
Remote-AP Operation	Configures when the virtual AP operates on a remote AP: <ul style="list-style-type: none"> • 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.

Table 27 Virtual AP Profile Parameters

Parameter	Description
Drop Broadcast and Multicast	<p>Select the Drop Broadcast and Multicast checkbox to filter out broadcast and multicast traffic in the air.</p> <p>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.</p> <p>IMPORTANT: If you enable this option, you must also enable the Broadcast-Filter ARP parameter in the stateful firewall configuration to prevent ARP requests from being dropped. To enable this setting:</p> <ol style="list-style-type: none"> 1. Navigate to Configuration > Stateful Firewall. 2. Click the Global Setting tab. 3. Select the Broadcast-Filter ARP checkbox. 4. Click Apply to save your settings before you return to the Virtual AP Profile. <p>Note also that although a virtual AP profile can be replicated from a master controller to local controllers, stateful firewall settings do not. If you select the Drop Broadcast and Multicast option for a Virtual AP Profile on a master controller, you must enable the Broadcast-Filter ARP setting on each individual local controller.</p>
Convert Broadcast ARP requests to unicast	<p>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 parameter is disabled by default.</p> <p>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 convert ARP requests directed to the broadcast address into unicast.</p> <p>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.</p>
Band Steering	<p>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.</p> <p>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.</p> <p>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. You must enable the Local Probe Response parameter in the Wireless LAN SSID profile for the SSID that requires band steering to use the Band Steering feature.</p> <p>To enable the local probe response parameter:</p> <ol style="list-style-type: none"> 1. Select the SSID profile associated with the Virtual AP profile. 2. Click the SSID profile Advanced tab. 3. Select the Local Probe Response checkbox. 4. Click Apply to save your settings before you return to the Virtual AP profile.

In the CLI

```
wlan ssid-profile corpnet
  essid Corpnet
  opmode wpa2-aes
wlan virtual-ap corpnet
  vlan 1
  aaa-profile corpnet
  ssid-profile corpnet
ap-group default
  virtual-ap corpnet
```

Guest WLAN

To configure Guest WLAN, the following basic steps are required.

- Configure the VLAN for guest users.
- Configure the guest role which only allows HTTP and HTTPS traffic from 9:00 a.m. to 5 p.m. on weekdays.
- Create and configure the virtual AP profile **guest** for the AP named “building3-lobby”:
 - Create a new virtual AP profile **guest**.
 - Select the predefined AAA profile **default-open**.
 - Create a new SSID profile **guest** to configure “Guest” for the SSID name and open system for the authentication.

The following sections describe how to do this using the WebUI and the CLI.

Configuring the 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.

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

```
vlan 2
interface vlan 2
  ip address <address> <netmask>
```

Configuring the 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 **IPv4 Session**. Click **Add**.

Select Service, then select svc-http from the drop-down list. For Time Range, select the time range you previously configured. Select **Add**. Add another rule for svc-https. Click **Apply**.

4. 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**.
5. Click **Apply**.

In the CLI

```
time-range workhours periodic
  weekday 09:00 to 17:00
ip access-list session restricted
  any any svc-http permit time-range workhours
  any any svc-https permit time-range workhours
user-role guest
  session-acl restricted
```

Configuring the Guest Virtual AP

In this example, you apply the **guest** virtual AP profile to a specific AP.



Note: Dell recommends that you 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** for the AP to display the profiles that you can configure for the AP.



Note: Selecting Wireless LAN allows you to exclude certain virtual AP profiles from being applied to this AP.

4. Select Virtual AP.
 - a. For Add a profile, select **NEW**.
 - b. Enter **guest**, and click **Add**.
 - c. Click **Apply**.
5. 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**.
6. 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**.
7. 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

```
wlan ssid-profile guest
  opmode opensystem
wlan virtual-ap guest
  vap-enable
  vlan 2
  deny-time-range workhours
  ssid-profile guest
  aaa-profile default-open
ap-name building3-lobby
  virtual-ap guest
```

Configuring High-throughput on Virtual APs

With the implementation of the IEEE 802.11n standard, high-throughput can be configured to operate on the 5 GHz and/or 2.4 GHz frequency band.

For high-throughput to function on a virtual AP profile for the assigned AP group or specific AP, high-throughput must be enabled within the assigned ht-ssid-profile and the radio-profile(s) for the desired frequency band(s).

By default, high-throughput is enabled; however, the examples in this section guide you through manually creating profiles and enabling high-throughput on the 5 GHz and 2.4 GHz frequency bands to ensure proper functionality of a virtual AP profile named “ht-vap-corpnet” assigned to an existing AP group named “ht-corpnet-aps.”



Note: For an example of 20 MHz channel versus 40 MHz channel pair configuration, see “20 MHz and 40 MHz Static Channel Assignments” on page 157.

This example includes the following tasks:

- Create two high-throughput radio profiles named “ht-radioa-corpnet” and “ht-radiog-corpnet.”
- Create and configure a 5 GHz radio profile named “ht-corpnet-a” and assign the high-throughput radio profile named “ht-radioa-corpnet.”
- Create and configure a 2.4 GHz radio profile named “ht-corpnet-g” and assign the high-throughput radio profile named “ht-radiog-corpnet.”
- Create and configure a high-throughput SSID profile named “ht-ssid-corpnet.”
- Create an SSID profile named “ht-corpnet” and assign the high-throughput SSID profile named “ht-ssid-corpnet.”
- Create a virtual AP profile named “ht-vap-corpnet” and assign the SSID profile named “ht-corpnet.”
- Assign the required profiles to an existing AP group named “ht-corpnet-aps.”

The following procedures are presented for the WebUI and the CLI.

In the WebUI

1. Navigate to **Configuration > Wireless > AP Configuration > AP Group** page.
2. Click **Edit** for the AP group ht-corpnet-aps.
3. Under the Profiles list, select **RF Management** to display the radio profiles.
4. Select the **802.11a radio profile**.



Note: This radio profile represents activity on the 5 GHz frequency band. Since the high-throughput IEEE 802.11n standard operates on the 5 GHz and/or 2.4 GHz frequency band, high-throughput can be enabled on 802.11a or 802.11g radio profiles.

- a. Select **New** from the 802.11a radio profile drop-down menu.
 - b. Enter **ht-corpnet-a** for the 802.11a radio profile name.
 - c. Select (check) the **High Throughput enable (radio)** checkbox to enable high-throughput. By default, this is enabled (checked).
 - d. Click **Apply**.
5. Select the **High-throughput Radio Profile** under the 802.11a radio profile.
 - a. Select **New** from the **High-throughput Radio Profile** drop-down menu.
 - b. Enter **ht-radioa-corpnet** for the high-throughput radio profile name.
 - c. Configure the high-throughput radio settings (see [Table 28](#) for details) and click **Apply**.

Table 28 *High-Throughput Radio Profile Configuration Parameters*

Parameter	Description
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.
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
Legacy station workaround	Select this option to enable interoperability for misbehaving legacy stations. This option is disabled by default, and should only be enabled under the supervision of Dell technical support.

6. Select the **802.11g** radio profile.



Note: This radio profile represents activity on the 2.4 GHz frequency band. Since the high-throughput IEEE 802.11n standard operates on the 5 GHz and/or 2.4 GHz frequency band, high-throughput can be enabled on 802.11a or 802.11g radio profiles.

- a. Select **New** from the 802.11g radio profile drop-down menu.
 - b. Enter **ht-corpnet-g** for the 802.11a radio profile name.
 - c. Select (check) the **High Throughput enable (radio)** checkbox to enable high-throughput. By default, this is enabled (checked).
 - d. Click **Apply**.
7. Select the **High-throughput Radio Profile** under the 802.11g radio profile.
 - a. Select **New** from the **High-throughput Radio Profile** drop-down menu.
 - b. Enter **ht-radiog-corpnet** for the high-throughput radio profile name.
 - c. Configure the high-throughput radio settings (see [Table 28](#) for details) and Click **Apply**.
 8. Select **Wireless LAN**, under the Profiles list, to reveal the WLAN profiles.
 9. Select the **Virtual AP** profile.
 - a. Select **New** from the **Add a Profile** drop-down menu.
 - b. Enter **ht-vap-corpnet** for the virtual AP profile name.
 - c. Click **Add**.
 - d. Select **New** from the **SSID Profile** drop-down menu associated with the “ht-vap-corpnet” virtual AP profile. The SSID Profile dialog box appears.
 - e. Enter **ht-corpnet** for the SSID profile name.

- f. Click **Apply** to create the SSID profile and return to the virtual AP profile page.
 - g. Click **Apply** on the virtual AP profile page.
10. Select the **ht-vap-corpnet** virtual AP profile.
 - a. Select **all** from the **Allowed band** drop-down menu.
 - b. Click **Apply**.
 11. Select the SSID profile **ht-corpnet**. The High-throughput SSID profile option will appear.
 12. Select the **High-throughput SSID Profile**.
 - a. Select **New** from the **High-throughput SSID Profile** drop-down menu.
 - b. Enter **ht-ssid-corpnet** for the high-throughput SSID profile name.
 - c. Click **Apply** to create the high-throughput SSID profile and assign it to the SSID profile.

In the CLI

```

rf ht-radio-profile ht-radioa-corpnet
rf ht-radio-profile ht-radiog-corpnet
rf dot11a-radio-profile ht-corpnet-a
  high-throughput-enable
  ht-radio-profile ht-radioa-corpnet
rf dot11g-radio-profile ht-corpnet-g
  high-throughput-enable
  ht-radio-profile ht-radiog-corpnet
wlan ht-ssid-profile ht-ssid-corpnet
  high-throughput-enable
wlan ssid-profile ht-corpnet
  ht-ssid-profile ht-ssid-corpnet
wlan virtual-ap ht-vap-corpnet
  allowed-bands all
  ssid-profile ht-corpnet
ap-group ht-corpnet-ap
  dot11a-radio-profile ht-corpnet-a
  dot11g-radio-profile ht-corpnet-g
  virtual-ap ht-vap-corpnet

```

Managing High-throughput Radio Profiles

Use the following command to create a high-throughput radio profile or edit an existing profile. For details see [Table 28](#).

```

rf ht-radio-profile <profile>
  40MHz-intolerance
  clone <profile>
  honor-40MHz-intolerance
  no
  single-chain-legacy

```

Advanced Configuration Options

This section describes advanced configuration options for APs and additional deployment information.

Configuring 802.11k Protocol

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 following procedure outlines 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.



Note: You cannot use spaces in profile names.

4. Configure your 802.11k radio settings. [Table 29](#) outlines the parameters you can configure in the 802.11k profile. Click **Apply** to save your settings.

Table 29 802.11k 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 <i>on-hook</i> 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 fulfil 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: <ul style="list-style-type: none">• 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

In the CLI

Use the following command to configure 802.11k profiles. The available parameters for this profile are described in [Table 29](#).

```
wlan dot11k <profile>
  bcn-measurement-mode {active|beacon-table|passive}
  clone <profile>
  dot11k-enable
  force-disassoc
```

Configuring RF Optimization

ArubaOS 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.

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 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 30](#) outlines the parameters. Click **Apply** to save your settings.

Table 30 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
Detect Association Failure	Enables or disables detection of station association failures. Default: Disabled
Detect interference	Select this checkbox to enable the interference detection. Default: Disabled
Interference Threshold	Percentage increase in the frame retry rate or frame receive error rate before interference monitoring begins on a given channel.
Interference Threshold Exceed Time	Amount of time the frame retry rate or frame receive error rate should be exceed by the threshold before interference is reported. Max 360000.
Interference Baseline Time	Time, in seconds, the air monitor should learn the state of the link between the AP and client to create frame retry rate (FRR) and frame receive error rate (FRER) baselines.
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

Table 30 RF Optimization Profile Parameters

Parameter	Description
Low RSSI Threshold	Minimum RSSI above which de-authorization messages should never be sent.
RSSI Check Frequency	Interval, in seconds, to sample RSSI.

In the CLI

Use the following command to configure RF Optimization profiles. The parameters described in [Table 30](#).

```
rf optimization-profile <profile>
clone <profile>
detect-association-failure
detect-interference
handoff-assist
interference-baseline <seconds>
interference-exceed-threshold <seconds>
interference-threshold <percent>
low-rssi-threshold <number>
no ...
rssi-check-frequency <number>
rssi-falloff-wait-time <seconds>
```

Configuring RF Event Configuration

The 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.



Note: 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.

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 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.

- Configure your settings as detailed in [Table 31](#) and click **Apply** to save your settings.

Table 31 RF Event 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.
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%.

In the CLI

Use the following command to configure RF event profiles. The available parameters for this profile are detailed in [Table 31](#).

```
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>
```

Changing AP Installation Modes

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.

In 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).

In 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

Configuring 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 re-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.



Note: Clients must support CSA in order to track the channel change without experiencing disruption.

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. 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**.

In the CLI

```
rf radio-profile <profile>  
  csa  
  csa-count <number>
```

20 MHz and 40 MHz Static Channel Assignments

With the implementation of the high-throughput IEEE 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 following channel configurations are now available in ArubaOS:

- A 20 MHz channel assignment consists of a single 20 MHz channel assignment. 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.

If high-throughput is disabled, a 40 MHz channel assignment can be configured, but only the primary channel assignment is utilized. The 20 MHz clients can also associate using this configuration, but only the primary channel is utilized.

Table 32 20 MHz and 40 MHz Static Channel Configuration Options

WebUI	CLI	Definition
Channel Text Field None Radio Button	<code>channel <num></code>	Entering a channel number in the CLI, or entering a channel number in the WebUI and selecting the None radio button, disables 40 MHz mode and activates 20 MHz mode for the entered channel.
Channel Text Field Above Radio Button	<code>channel <num>+</code>	Entering a channel number with a plus (+) sign in the CLI, or entering a channel number and selecting the Above radio button in the WebUI, selects a primary and secondary channel for 40 MHz mode. The number entered becomes the primary channel and the secondary channel is determined by increasing the primary channel number by 4. Example: 157+ represents 157 as the primary channel and 161 as the secondary channel.
Channel Text Field Below Radio Button	<code>channel <num>-</code>	Entering a channel number with a minus (-) sign in the CLI, or entering a channel number and selecting the Below radio button in the WebUI, selects a primary and secondary channel for 40 MHz mode. The number entered becomes the primary channel and the secondary channel is determined by decreasing the primary channel number by 4. Example: 157- represents 157 as the primary channel and 153 as the secondary channel.

The example in this section illustrates a static channel assignment and assumes that the radio and regulatory domain profiles being configured were previously created and assigned to an existing AP group named “ht-corpnet-ap.” These settings also allow for the default ARM profile settings, see “Automatic Channel and Transmit Power Selection Using ARM” on page 159, and Dell’s recommended high-throughput channel assignments for the 802.11a and 802.11b/g bands:

1. Enter a valid country code (US) for the “default” regulatory domain profile. This will determine the available channels.
2. Configure a 40 MHz channel (bonded pair) for an 802.11a (5 GHz) radio profile named “ht-corpnet-a.”
3. Configure a 20 MHz channel for an 802.11g (2.4 GHz) radio profile named “ht-corpnet-g.”



Note: If you want the channel assignments to utilize high-throughput, ensure that high-throughput is enabled within the radio profile. For details, see “Configuring High-throughput on Virtual APs” on page 125.

Configuring Channels in the WebUI

1. Navigate to **Configuration > Wireless > AP Configuration > AP Group** page.
2. Click **Edit** for the AP group ht-corpnet-ap.
3. Under the Profiles list, select **AP** to display the AP profiles.
4. Select the **Regulatory Domain** profile named “default.”
5. Select **US - United States** from the **Country Code** drop-down menu.
6. Click **Apply**.
7. Under the Profiles list, select **RF Management** to display the radio profiles.
8. Select the **802.11a** radio profile named “ht-corpnet-a.”
9. 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.
10. Click **Apply**.
11. Select the **802.11g** radio profile named “ht-corpnet-g.”

12. Enter **1** in the **Channel** text field and select the **None** radio button. In this instance, channel 1 is the assigned 20 MHz channel and 40 MHz mode is disabled and click **Apply**.

Configuring channels in the CLI

```
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
```

Automatic Channel and Transmit Power Selection Using ARM

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 [Chapter 5, “Adaptive Radio Management \(ARM\)”](#) on page 139.

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.

Best Practices

When deploying APs across low-speed or high-latency links, Dell recommends the following:

- Connect APs and controllers over a link with a capacity of 1 Mbps or greater.
- Maintain a minimum link speed of 64 Kbps per GRE tunnel 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 [Chapter 4, “Access Points”](#).
- 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 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**, then **AP system profile**. The configuration settings are displayed in **Profile Details**.
4. Under **Profile Details**:
 - a. At the **Bootstrap threshold**, enter 30.
 - b. Click **Apply**.

Configuring in the CLI

```
ap system-profile <profile>
  bootstrap-threshold 30
```

Prioritizing AP heartbeats 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**, 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**.

Prioritizing AP heartbeats in the CLI

```
ap system-profile <profile>
  heartbeat-dscp <number>
```

AP Redundancy

In conjunction with the controller redundancy features described in [Chapter 20, “VRRP”](#) 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 [Chapter 6, “Remote Access Points”](#).

AP failback

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.

To configure this feature you must:

- Configure the LMS IP address
- Configure the backup LMS IP address
- Enable LMS preemption
- Configure the LMS hold-down timer

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.

Configuring 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.
 - c. Click (select) **LMS Preemption**. This is disabled by default.
6. Click **Apply**.

Configuring in 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.

Configuring 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, 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**.

Configuring in 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>|ssid <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>|ssid <name>|ip-addr <ipaddr>}
show ap database
show ap details {ap-name <name>|bssid <name>|ip-addr <ipaddr>}
```

Managing 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.

Disabling LEDs in the WebUI

An AP system profile's LED operating mode affects LEDs on all APs using that profile.



Note: This option is only available on the W-AP120 Series.

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 (W-API20 series only)** parameter.
5. From the drop-down list, select **off**.
6. Click **Apply**.

Enable or Disable LEDs in 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}
```

Configuring Blinking LEDs in the CLI

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 @@@@
```

This chapter outlines how to configure the ARM function to automatically select the best channel and transmission power settings for each AP on your WLAN. This chapter describes the following topics:

- “ARM Overview” on page 139
- “ARM Profiles” on page 140
- “Multi-Band ARM and 802.11a/802.11g Traffic” on page 147
- “Band Steering” on page 147
- “Traffic Shaping” on page 148
- “Spectrum Load Balancing” on page 149
- “RX Sensitivity Tuning Based Channel Reuse” on page 149
- “Non-802.11 Noise Interference Immunity” on page 150
- “ARM Metrics” on page 150
- “ARM Troubleshooting” on page 151

ARM Overview

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 DellAP 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 supports APs with the 802.11n standard, ensuring seamless integration of 802.11n devices into your RF domain. 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, your AP will dynamically scan all 802.11 channels within its 802.11 regulatory domain at regular intervals and will report everything it sees to the controller on each channel it scans. 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 Metrics” on page 150.](#))

Noise and Error Monitoring

An AP configured with ARM is aware of both 802.11 and non-802.11 noise, and will adjust to a better channel if it reaches a configured threshold for either noise, MAC errors or PHY errors. 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.

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.

ARM Profiles

You configure ARM by defining ARM *profiles*, a set of configuration parameters that you can apply as needed to an AP group or to individual APs. Dell controllers have one preconfigured ARM profile, called **default**. Most network administrators will find that this one default ARM profile is sufficient to manage all the Dell APs on their WLAN. Others may want to define multiple profiles to suit their APs' varying needs.

When managing ARM profiles, you should first consider whether or not all the APs on your WLAN operate in similar environments and manage similar traffic loads and client types.

If your APs' environment and traffic loads are mostly the same, you can use the default ARM profile to manage all the APs on your WLAN. If you ever modify the default profile, all APs on the WLAN will be updated with the new settings. If, however, you have APs on your WLAN that are in different physical environments, or your APs each manage widely varying client loads or traffic types, you should consider defining additional ARM profiles for your AP groups. [Table 33](#) describes different WLAN environments, and the type of ARM profiles appropriate for each.

Table 33 ARM Profile Types

ARM Profiles	Example WLAN Description
default profile only	<ul style="list-style-type: none"> • A warehouse where the physical environment is nearly the same for all APs, and each AP manages the same number of clients and traffic load. • A training room, where the clients are evenly spaced throughout the room, have the same security requirements and are using the same amount of network resources.
multiple profiles	<ul style="list-style-type: none"> • Universities where APs are in different building types (open auditoriums, small brick classrooms), some APs must support VoIP or video streaming, and mobile clients are constantly moving from one AP coverage area to another. • Healthcare environments where some APs must balance the network demands of large digital radiology files, secure electronic patient record transfers, diagnostic videos, and collaborative VoIP sessions, while other APs (like those in a lobby or cafeteria) support only lower-priority traffic like Internet browsing.

You assign ARM profiles to AP groups by associating an ARM profile with that AP group's 802.11a or 802.11g RF management profile. For details on associating an ARM profile with an AP group, see [“Assigning an ARM Profile to an AP Group” on page 146](#).

There are two ways to create a new ARM profile. You can make an entirely new profile with all default settings, or you can create a new profile based upon the settings of an existing profile by making a copy of that other profile.

Creating a New ARM Profile

To create a new ARM profile with all default settings via the WebUI:

1. Select **Configuration > 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**.

To create a new ARM profile via the command-line interface, access the CLI in config mode and issue the following command.

```
rf arm-profile <profile>
```

where <profile> is a unique name for the new ARM 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

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.

```
rf arm-profile <newprofile> clone <profile>
```

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. 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

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 using the WebUI:

1. Select **Configuration > 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 command

```
no rf arm-profile <profile>
```

where <profile> is the name of the ARM profile you wish to remove.

Configuring ARM Settings via the WebUI

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.

To change an ARM profile:

1. Select **Configuration > 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 edit. The **Adaptive Radio Management (ARM) profile** window opens.
5. Change any of the ARM settings described in the table below, then click **Apply** to save your changes.

Table 34 ARM Profile Configuration Parameters

Setting	Description
Assignment	<p>Activates one of four ARM channel/power assignment modes.</p> <ul style="list-style-type: none"> ● 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. <p>Default: single-band</p>

Table 34 ARM Profile Configuration Parameters

Setting	Description
Client Aware	<p>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.)</p> <p>If Client Aware is disabled, the AP may change to a more optimal channel, but this change may also disrupt current client traffic.</p> <p>Default: enabled</p>
Min Tx EIRP	<p>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 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.</p> <p>Default: 9 dBm</p> <p>NOTE: 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.</p>
Max Tx EIRP	<p>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.</p> <p>Default: 127 dBm</p> <p>NOTE: Power settings will not change if the Assignment option is set to disabled or maintain.</p>
Multi Band Scan	<p>If enabled, single radio channel APs scans for rogue APs across multiple channels. This option requires that Scanning is also enabled.</p> <p>(The Multi Band Scan option does not apply to APs that have two radios, such as an Dell AP-65 or AP-70, 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.)</p> <p>Default: disabled</p>
Rogue AP Aware	<p>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.</p> <p>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.</p> <p>Default: disabled</p>
Scan Interval	<p>If Scanning is enabled, the Scan Interval defines how often the AP will leave its current channel to scan other channels in the band.</p> <p>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.</p> <p>The supported range for this setting is 0–2,147,483,647 seconds.</p> <p>Default: 10 seconds</p>
Active Scan	<p>When the Active Scan checkbox is selected, 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 <i>not be enabled</i> except under the direct supervision of Dell Support.</p> <p>Default: disabled</p>

Table 34 ARM Profile Configuration Parameters

Setting	Description
Scanning	<p>The Scanning checkbox enables or disables AP scanning across multiple channels. Disabling this option also disables the following scanning features:</p> <ul style="list-style-type: none"> • Multi Band Scan • Rogue AP Aware • Voip Aware Scan • Power Save Scan <p>Do not disable Scanning unless you want to disable ARM and manually configure AP channel and transmission power. Default: enabled</p>
Scan Time	<p>The amount of time, in milliseconds, an AP will step out of the current channel to scan another channel. The supported range for this setting is 0–2,147,483,647 seconds. Dell recommends a scan time between 50–200 msec. Default: 110 msec</p>
VoIP Aware Scan	<p>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</p>
Power Save Aware Scan	<p>If enabled, the AP will not scan a different channel if it has one or more clients that is in power save mode. Default: enabled</p>
Ideal Coverage Index	<p>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. The range of possible values is 2–20. Default: 10 For additional information on how this the Coverage Index is calculated, see “ARM Metrics” on page 150</p>
Acceptable Coverage Index	<p>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. The range of possible values is 1–6. Default: 4</p>
Free Channel Index	<p>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 Metrics” on page 150</p>
Backoff Time	<p>After an AP changes channel or power settings, it waits for the backoff time interval before it asks for a new channel/power setting. The range of possible values is 120–3600 seconds. Default: 240 seconds</p>
Error Rate Threshold	<p>The minimum percentage of PHY errors and MAC errors in the channel that will trigger a channel change. Default: 50%</p>
Error Rate Wait Time	<p>Minimum time in seconds the error rate has to exceed the Error Rate Threshold before it triggers a channel change. Default: 30 seconds</p>
Noise Threshold	<p>Maximum level of noise in channel that triggers a channel change. The range of possible 0–2,147,483,647 dBm. Default 75 dBm</p>

Table 34 ARM Profile Configuration Parameters

Setting	Description
Noise Wait Time	Minimum time in seconds the noise level has to exceed the Noise Threshold before it triggers a channel change. The range of possible values is 120–3600 seconds. Default: 120 seconds
Minimum Scan Time	Minimum number of times a channel must be scanned before it is considered for assignment. The supported range for this setting is 0–2,147,483,647 scans. Dell recommends 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. The supported range for this setting is 0–20000000 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 (e.g. 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

Configuring ARM Settings via 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 34 on page 142](#). 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 ARM profile mode.

Use the following command to create or modify an ARM profile:

```
rf arm-profile <profile>
  40MHz-allowed-bands {All|None|a-only|g-only}
  acceptable-coverage-index <number>
  active-scan (not intended for use)
  assignment {disable|maintain|multi-band|single-band}
  backoff-time <seconds>
  client-aware
  clone <profile>
  error-rate-threshold <percent>
  error-rate-wait-time <seconds>
  free-channel-index <number>
  ideal-coverage-index <number>
  load-aware-scan-threshold <Mbps>
  max-tx-power <dBm>
  min-scan-time <# of scans>
  min-tx-power <dBm>
  mode-aware
  multi-band-scan
  no
  noise-threshold <number>
  noise-wait-time <seconds>
  ps-aware-scan
  rogue-ap-aware
```

```
scan-interval <seconds>
scan-time <milliseconds>
scanning
voip-aware-scan
```

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.

Assigning a Profile via the WebUI

To assign an ARM profile to an AP group via the Web User Interface:

1. Select **Configuration > AP Configuration**.
2. If it is not already selected, click the **AP Group** tab.
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 profile to an AP group's 802.11a radio profile, expand the **802.11a radio profile** section.
 - To assign a new 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 select an ARM profile for another profile.
9. Click **Apply** to save your changes.

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 > 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.

Assigning a Profile via 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:

```
rf dot11a-radio-profile <ap_profile>
   arm-profile <arm_profile>
```

and

```
rf dot11g-radio-profile <ap_profile>
   arm-profile <arm_profile>
```

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.

Multi-Band ARM and 802.11a/802.11g Traffic

Dell recommends using 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. 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 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.

Note: The Band Steering feature may not work correctly unless you use the enable the "Local Probe Response" parameter in the Wireless LAN SSID profile for the SSID that requires band steering. You can enable the local probe response parameter using the CLI command `wlan ssid-profile <profile> local-probe-response`, or via the WebUI by navigating to **Configuration>All Profiles**, expanding the **Wireless LAN** and **SSID Profile** menus, then selecting the **SSID profile** and checking the **Local Probe Response** checkbox in the **SSID Profile Details** window.



Enabling Band Steering via the WebUI

Band steering is configured in a virtual AP profile. Use the following procedure to enable or disable Band Steering using the WebUI:

1. Select **Configuration > 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** button. 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. Click **Apply** to save your changes.

Enabling Band Steering via 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.

```
wlan virtual-ap <profile> band-steering
```

To disable band steering, include the **no** parameter

```
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.

```
ap-group <name> virtual-ap <profile>  
ap-name <name> virtual-ap <profile>
```

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 via the WebUI

Traffic shaping is configured in a traffic management profile. To configure traffic shaping via the WebUI:

1. Select **Configuration > 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, then click **Add**. 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** to save your changes.

Enabling Traffic Shaping via the WebUI

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-access
```

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>
```

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 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.

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.

For details on modifying 802.11a or 802.11g RF management profiles, refer to [“RF Management \(802.11a and 802.11g\) Profiles” on page 204](#).

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.

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 \(802.11a and 802.11g\) Profiles” on page 204](#).

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 and 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.

You can manage Non-802.11 Noise Immunity settings through the 802.11g RF management profile. Do not raise the noise immunity feature’s default setting if the RX Sensitivity Tuning Based Channel Reuse feature is also enabled. A level-3 to level-5 Noise Immunity setting is not compatible with the Channel Reuse feature. For details refer to [“Mesh Radio Profiles” on page 199](#).

ARM 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 APs 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

ARM Troubleshooting

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 following troubleshooting tips.

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 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 are using 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 not disabled. 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 ARM noise checking is enabled. 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 Secure Remote Access Point Service allows AP users, at remote locations, to connect to a controller over the Internet. Since 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 APs. This chapter discusses the following topics:

- “Overview” on page 153
- “Configuring the Secure Remote Access Point Service” on page 154
- “Deploying a Branch Office/Home Office Solution” on page 162
- “Enabling Double Encryption” on page 166
- “Advanced Configuration Options” on page 167

Overview


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.

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 is used to secure AP-to-controller communication. (Dell recommends this deployment 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 23 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 24 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. (Dell recommends this deployment 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 25 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 26). 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 network by the master controller.

Figure 26 Remote AP in a Multi-Controller Environment



Configuring the Secure Remote Access Point Service

The tasks for configuring an 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.
- Configure the remote AP user role.
User roles and policies require the Policy Enforcement Firewall Next Generation (PEFNG) license. To configure and assign specific user roles you must install the Policy Enforcement Firewall Virtual Private Network (PEFV) license. The example in this section configures a custom user role and policy. You must

install the PEFNG and PEFV license in the controller, as described in [Chapter 27, “Software Licenses”](#) . . . Configure the authentication server that will validate the username and password for the remote AP.

- Provision the AP with IPSec settings, including the username and password for the AP, before you install it at the remote location.

ArubaOS supports multiple remote AP modes of operation. By default, the remote AP operates in standard mode. This mode enables the virtual AP when the remote AP connects to the controller. The information in this section assumes the default mode of operation. For information on remote AP modes of operation, refer to [“Advanced Configuration Options” on page 167](#).

Configure a Public IP Address for the Controller

The remote AP requires an IP address to which it can connect in order to establish a VPN tunnel to the controller. This can be either a routable IP address that 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

```
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 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 [Chapter 15, “Virtual Private Networks”](#) . 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 (check) **Enable L2TP**.
3. Make sure that only 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**.



Note: 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>
```

Configure the Remote AP User Role

Once the remote AP is authenticated for the VPN and established a IPsec connection, it is assigned a role. This role is a temporary role assigned to the AP until it completes the bootstrap process after which it inherits the ap-role. The appropriate ACLs need to be enabled to permit traffic from the controller to the AP and back to facilitate the bootstrap process.



Note: User roles and policies require the PEFNG license. You must install the PEFNG license, as described in [Chapter 27, "Software Licenses"](#).

To configure the user role, you create a policy that permits the following traffic:

- AP control traffic via the Dell PAPI protocol
- GRE tunnel traffic
- Layer-2 Tunneling Protocol (L2TP) traffic
- TFTP traffic from the remote AP to the controller
- FTP traffic from the remote AP to the controller

Then, you create a user role that contains this policy.

Using the WebUI

1. Navigate to the **Configuration > Security > Access Control > Policies** page.
2. Click **Add** to create a policy.
3. Enter the Policy Name (for example, remote-AP-access).
4. From the **Policy Type** drop-down list, select **IPv4 Session**.
5. To create the first rule:
 - a. Under **Rules**, click **Add**.
 - b. For **Source**, select **any**.
 - c. For **Destination**, select **any**.
 - d. For **Service**, select **service**, then select **svc-papi**.
 - e. Click **Add**.

6. To create the next rule:
 - a. Under Rules, click **Add**.
 - b. For Source, select **any**.
 - c. For Destination, select **any**.
 - d. For Service, select **service**, then select **svc-gre**.
 - e. Click **Add**.
7. To create the next rule:
 - a. Under Rules, click **Add**.
 - b. For Source, select **any**.
 - c. For Destination, select **any**.
 - d. For Service, select **service**, then select **svc-l2tp**.
 - e. Click **Add**.
8. To create the next rule:
 - a. Under Rules, click **Add**.
 - b. For Source, select **any**.
 - c. For Destination, select **alias**, then select **mswitch**.
 - d. For Service, select **service**, then select **svc-tftp**.
 - e. Click **Add**.
9. To create the next rule:
 - a. Under Rules, click **Add**.
 - b. For Source, select **any**.
 - c. For Destination, select **alias**, then select **mswitch**.
 - d. For Service, select **service**, then select **svc-ftp**.
 - e. Click **Add**.
10. Click **Apply**.
11. Click the **User Roles** tab.
 - a. Click **Add**.
 - b. Enter the Role Name (for example, RemoteAP).
 - c. Click **Add** under Firewall Policies.
 - d. In the Choose from Configured Policies menu, select the policy you just created.
 - e. Click **Done**.
12. Click **Apply**.

Using CLI

```
ip access-list session <policy>
  any any svc-papi permit
  any any svc-gre permit
  any any svc-l2tp permit
  any alias mswitch svc-tftp permit
  any alias mswitch svc-ftp permit

user-role <role>
  session-acl <policy>
```

Configure VPN Authentication

Before you enable VPN authentication, you must configure the authentication server(s) and server group that the controller will use to validate the remote AP. When you provision the remote AP, you configure IPsec settings for the AP, including the username and password. This username and password must be validated by an authentication server before the remote AP is allowed to establish a VPN tunnel to the controller. The authentication server can be any type of server supported by the controller, including the controller's internal database.



Note: For security purposes, Dell best practices is to assign a unique username and password for each remote AP.

For more information about configuring authentication servers and server groups, refer to [Chapter 8, “Authentication Servers”](#).

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > L3 Authentication** page.
2. In the Profiles list, select the **VPN Authentication Profile > default-rap**.
3. For Default Role, enter the user role you created previously (for example, RemoteAP).



Note: User roles and policies require the PEFNG and PEFV license. You must install the PEFNG and PEFV license, as described in [Chapter 27, “Software Licenses”](#).

4. Click **Apply**.
5. In the Profile list, under VPN Authentication Profile, select **Server Group**.
6. Select the server group from the drop-down menu and click **Apply**.

Using CLI

```
aaa server-group <group>
  auth-server <server>
aaa authentication vpn default-rap
  default-role <role>
  server-group <group>
```

Configuring Internal Database for Authentication

You can use the controller's internal database as an authentication server. To configure the internal database for a remote AP user, do the following:

1. Configure a public IP address for the controller.
2. Configure the VPN server on the controller.
3. Configure the remote AP user role.
4. Configure VPN authentication using the internal database.
5. Add the user to the internal database.

The information in this section assumes you have configured a public IP address for the controller and the VPN server. For information about configuring the public IP address, see [“Configure a Public IP Address for the Controller” on page 155](#). For information about configuring the VPN server, see [“Configure the VPN Server” on page 155](#).

Using the WebUI

The following procedure illustrates the steps to configure an internal database for a remote AP user. To configure the user role, you first create a policy that permits the following traffic:

- AP control traffic via the Dell PAPI protocol
- GRE tunnel traffic
- ESP tunnel traffic
- Layer-2 Tunneling Protocol (L2TP) traffic
- TFTP traffic
- FTP traffic

Then, you create a user role that contains this policy.

1. Navigate to the **Configuration > Security > Access Control > Policies** page.
2. Click **Add** to create a policy.
3. Enter the Policy Name (for example, rap_policy).
4. From the **Policy Type** drop-down list, select **IPv4 Session**.
5. To create the first rule:
 - f. Under Rules, click **Add**.
 - g. For Source, select **any**.
 - h. For Destination, select **any**.
 - i. For Service, select **service**, then select **svc-papi**.
 - j. Click **Add**.
6. To create the next rule:
 - a. Under Rules, click **Add**.
 - b. For Source, select **any**.
 - c. For Destination, select **any**.
 - d. For Service, select **service**, then select **svc-l2tp**.
 - e. Click **Add**.
7. To create the next rule:
 - a. Under Rules, click **Add**.
 - b. For Source, select **any**.
 - c. For Destination, select **any**.
 - d. For Service, select **service**, then select **svc-gre**.
 - e. Click **Add**.
8. To create the next rule:
 - a. Under Rules, click **Add**.
 - b. For Source, select **any**.
 - c. For Destination, select **any**.
 - d. For Service, select **service**, then select **svc-esp**.
 - e. Click **Add**.
9. To create the next rule:
 - a. Under Rules, click **Add**.
 - b. For Source, select **any**.

- c. For Destination, select **any**.
 - d. For Service, select **service**, then select **svc-tftp**.
 - e. Click **Add**.
10. To create the next rule:
- a. Under Rules, click **Add**.
 - b. For Source, select **any**.
 - c. For Destination, select **any**.
 - d. For Service, select **service**, then select **svc-ftp**.
 - e. Click **Add**.
11. Click **Apply**.
12. Click the **User Roles** tab.
- a. Click **Add**.
 - b. Enter the Role Name (for example, `rap_role`).
 - c. Click **Add** under Firewall Policies.
 - d. In the Choose from Configured Policies menu, select the policy you just created.
 - e. Click **Done**.
13. Click **Apply**.

Configure VPN authentication using the internal database

1. Navigate to the **Configuration > Security > Authentication > L3 Authentication** page.
2. In the Profiles list, select **VPN Authentication Profile**.
3. For Default Role, enter the user role you created previously (for example, `rap_role`).
4. Click **Apply**.
5. In the Profile list, under VPN Authentication Profile, select **Server Group**.
6. Select the **internal** server group from the drop-down menu.
7. Click **Apply**.

Add the user to the internal database

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 user name and password.
5. Click **Enabled** to activate this entry on creation.
6. Click **Apply** to apply the configuration. Note that the configuration does not take effect until you perform this step.
7. At the **Servers** page, click **Apply**.

Using CLI to configure the internal DB for a RAP user

```
ip access-list session rap_policy
  any any svc-papi permit
  any any svc-l2tp permit
  any any svc-gre permit
  any any svc-esp permit
  any any svc-tftp permit
  any any svc-ftp permit
```

```
user-role rap_role
session-acl rap_policy
```

Configure VPN authentication using the internal database:

```
aaa authentication vpn
default-role rap_role
server-group internal
```

Add the user to the internal database:

```
local-userdb add username rapuser1 password <password>
```

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. See [Appendix G, “Provisioning RAP at Home”](#) for more information about provisioning remote AP at home.

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.



Note: 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



Note: 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**.

Creating a Remote AP Whitelist

Remote AP whitelist is the list of approved AP's 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 the **New** button and provide the following details:
 - **AP MAC Address**—Mandatory parameter. Enter the MAC address of the AP.
 - **Username**—Enter a username that will be used when the AP is provisioned.
 - **AP Group**—Select a group to add the AP.
 - **AP Name**—Enter a name for the AP. If an AP name is not entered, 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 the **Add** button to add the remote AP to the whitelist.

Revoking an AP

In some cases, if an AP in the whitelist is retired from active usage, you can set the AP as revoked. This option restricts the AP from connecting to your controller. To revoke a remote AP:


1. Select an AP from the whitelist by selecting the checkbox.
2. Click the **Modify** button.
3. Select the checkbox under the **Revoked** column.
4. Click the **Update** button.

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 like printers and servers but traffic to and from these resources must not impact the corporate head office.

The [Figure 27](#) 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 27 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.

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 (by default, the remote AP operates in standard mode)
- 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



Note: 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 AP and to 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.

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 35 Rap Console Summary Tab Information

Summary Table Name	Basic View Information	Advanced View Information
Wired Ports Status	<ul style="list-style-type: none"> ● Port: Port numbers of the wired ports on the AP. ● Status: Current status of each port (<i>Connected, Link Down</i> or <i>Disabled</i>). 	<p>The advanced view of the Wired Access Ports table displays the following data:</p> <ul style="list-style-type: none"> ● Port: Port numbers of the wired ports on the AP. ● Status: Current status of each port (<i>Connected, Link Down</i> or <i>Disabled</i>). ● 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: <i>Bridge, Tunnel</i> or <i>Split Tunnel</i>. ● Users: Number 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	<ul style="list-style-type: none"> ● SSID: Name of the SSID. ● Status: SSID Status (up, down, or disabled). ● Band: Radio band available on the SSID. 	<ul style="list-style-type: none"> ● 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 (<i>Bridge, Tunnel</i> or <i>Split-Tunnel</i>). ● EIRP: Equivalent Isotropic Radiated Power, in dBm. ● Noise floor: The 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	<ul style="list-style-type: none"> ● MAC Address: MAC address of the wired user. ● IP address: IP address of the wired user. 	<ul style="list-style-type: none"> ● MAC Address: MAC address of the wired user. ● IP address: IP address of the wired user. ● Port: AP port used by the wired user.

Table 35 Rap Console Summary Tab Information

Summary Table Name	Basic View Information	Advanced View Information
Wireless User	<ul style="list-style-type: none"> ● MAC Address: MAC address of the wireless user. ● IP address: IP address of the wireless user. 	<ul style="list-style-type: none"> ● 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 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: The Receive Signal Strength Indicator (RSSI) value displayed in the output of this command represents signal strength as a signal to noise ratio.
Device Info	<ul style="list-style-type: none"> ● 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. ● Ims: IP address of the local controller. 	N/A
Uplink Info	<p>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.</p> <p>Active uplink information, including:</p> <ul style="list-style-type: none"> ● Interface name ● Port speed ● IP address <p>Standby link information, including:</p> <ul style="list-style-type: none"> ● Name (3G) ● Device connected (yes/no) ● Provisioned (yes/no) ● IP address ● Device ● User ● Password 	N/A

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 36](#).

Table 36 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 that was 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 that was 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. You can also use the **Reboot AP Now** button at the bottom of the Diagnostic window reboots 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 Double Encryption

The double encryption feature applies only for traffic to and from a wireless client that is connected to a tunneled SSID. When this feature is enabled, all traffic (which is already encrypted using Layer-2 encryption) is re-encrypted in the IPSec tunnel. When this feature is 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.

Using the WebUI

1. Navigate to the **Configuration > Wireless > AP Configuration > AP Specific** page. Click **Edit** for the remote AP.
2. Under **Profiles**, select **AP**, then select AP system profile.
3. Under **Profile Details**, select the AP system profile for this AP from the drop-down menu. Select **Double Encrypt**. Click **Apply**.

Using CLI

```
ap system-profile <profile>
    double-encrypt
ap-name <name>
    ap-system-profile <profile>
```



Note: Dell recommends that double-encryption not be turned on for inter-device communication over untrusted networks, as doing so is redundant and adds significant processing overhead for APs.

Advanced Configuration Options

This section describes the following features designed to enhance your remote AP configuration:

- [“Understanding Remote AP Modes of Operation” on page 167](#)
- [“Fallback Mode” on page 169](#)
- [“DNS Controller Setting” on page 177](#)
- [“Backup Controller List” on page 177](#)
- [“Remote AP Failback” on page 178](#)
- [“Access Control Lists and Firewall Policies” on page 1801](#)
- [“Split Tunneling” on page 180](#)
- [“Wi-Fi Multimedia” on page 186](#)



Note: The information in this section assumes you have already configured the remote AP functionality, as described [“Configuring the Secure Remote Access Point Service” on page 154](#).

Understanding Remote AP Modes of Operation

[Table 37](#) 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.



Note: 802.1x and PSK authentication is supported when you configure the remote AP to operate in bridge or split-tunnel mode.

Table 37 Remote AP Modes of Operation and Behavior

Remote AP Operation Setting					
	all	bridge	split-tunnel	tunnel	decrypt-tunnel
all		<p>Management frames on AP.</p> <p>Frames are bridged between wired and wireless interfaces.</p> <p>No frames are tunneled to the controller.</p> <p>Station acquires its IP address locally from an external DHCP server.</p>	<p>Management frames on AP.</p> <p>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.</p> <p>Typically, the station obtains an IP address from a VLAN on the controller.</p> <p>Typically, the AP has ACLs that forward corporate traffic through the tunnel and source NAT the non-corporate traffic to the Internet.</p>	<p>Management frames as per local-probe response and association on APs.</p> <p>Frames are GRE tunneled to the controller to an untrusted tunnel.</p> <p>100% of station frames are tunneled to the controller.</p>	<p>Management frames on AP.</p> <p>Frames are always GRE tunneled to controller.</p>
always	<p>ESSID is always up when the AP is up regardless if the controller is reachable.</p> <p>Supports PSK ESSID only.</p> <p>SSID configuration stored in flash on AP.</p>	<p>Provides an SSID that is always available for local access.</p>	Not supported	Not supported	Not supported
	all	bridge	split-tunnel	tunnel	
backup	<p>ESSID is only up when controller is unreachable.</p> <p>Supports PSK ESSID only.</p> <p>SSID configuration stored in flash on AP.</p>	<p>Provides a backup SSID for local access only when the controller is unreachable.</p>	Not supported	Not supported	Not supported
persistent	<p>ESSID is up when the AP contacts the controller and stays up if connectivity is disrupted with the controller.</p> <p>SSID configuration obtained from the controller.</p> <p>Designed for 802.1x SSIDs.</p>	<p>Same behavior as standard, described below, except the ESSID is up if connectivity to the controller is lost.</p>	Not supported	Not supported	Not supported

Table 37 Remote AP Modes of Operation and Behavior

Remote AP Operation Setting	Forward Mode Setting				
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

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 becomes 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 37](#) 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.

This section describes the following topics:

- [“Configuring the fallback mode” on page 169](#)
- [“Configuring the DHCP Server on the Remote AP” on page 171](#)
- [“Advanced Backup Configuration Options” on page 1733](#)

Configuring the fallback mode

To configure the fallback mode, you must

- Configure the AAA profile.
- Configure the virtual AP profile

Using WebUI to configure the AAA profile

The AAA profile defines the authentication method and the default user role for unauthenticated users.



Note: 802.1x and PSK authentication is supported when configuring bridge or split tunnel mode.

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**.



Note: 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.

Using CLI

```
aaa profile <name>
  initial-role <role>
  authentication-dot1x <dot1x-profile>
  dot1x-default-role <role>
  dot1x-server-group <group>
```

Using the WebUI to configure virtual AP profile

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 [“Configuring Profiles” on page 105 in Chapter 4, “Access Points”](#).

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**.

Using 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. By default, the AP assigns IP addresses from the DHCP pool 192.168.11.0/24, 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. By default, the AP uses IP address 192.168.11.1 for the DHCP server, the DHCP router and the DHCP DNS server.
- Enter the amount of days the assigned IP address is valid (also known as the remote AP DHCP lease). By default, the lease does not expire, 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.



Note: 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 the fallback mode” on page 169](#)).

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

```
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-dhcp-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>
```

Advanced Backup Configuration 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.



Note: 802.1x and PSK authentication is supported when configuring bridge or split tunnel mode.

- 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 the fallback mode” on page 169](#).

- 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 171](#).

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.



Note: The remote AP 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.

Using the WebUI to configure the session ACL

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**.



Note: 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**.

Using the WebUI to configure the AAA profile

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**.



Note: 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 your remote AP configuration, then click **Apply**.

Using the WebUI to define the backup configuration

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**.

Using the CLI to configure the session ACL

```
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 route src-nat**.

```
user-role <role>
  session-acl <policy>
```

Using the CLI to configure the AAA profile

```
aaa profile <name>
  initial-role <role>
```

You can define other parameters as needed.

Using the CLI to define the backup configuration

```
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>
```

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 setup 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. Dell recommends using 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 177](#).

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.



Note: Reprovisioning the AP causes it to automatically reboot.

Specify the DNS name using 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 161](#).


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.

If the remote AP loses connectivity on the IPsec tunnel to the controller, the remote AP establishes connectivity with a backup controller from the list and automatically reboots. Network connectivity is lost during this time. As described in the section [“Remote AP Failback” on page 178](#), 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 28 Sample Backup Controller Scenario



arun_023

Configuring the LMS and backup LMS IP addresses using 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 using CLI

```
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>
```

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.

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. Click (select) **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**.

Using the CLI

```
ap system-profile <profile>
  lms-preemption
  lms-hold-down period <seconds>
```

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 CLI or the WebUI to enable the local network access.

Using 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 29 Enable Remote AP Local Network Access

Session ACL	ap-uplink-acl	Corporate DNS Domain	<input type="text"/> Delete Add
Maintenance Mode	<input type="checkbox"/>	WISPr Location-ID ISO Country Code	<input type="text"/>
WISPr Location-ID E.164 Country Code	<input type="text"/>	WISPr Location-ID E.164 Area Code	<input type="text"/>
WISPr Location-ID SSID/Zone	<input type="text"/>	WISPr Operator Name	<input type="text"/>
WISPr Location Name	<input type="text"/>	Remote-AP Local Network Access	<input checked="" type="checkbox"/>

5. Click **Apply**.

Using CLI

- To enable, enter:


```
ap system-profile <ap-profile> rap-local-network-access
```
- To disable, enter:


```
ap system-profile <ap-profile> no rap-local-network-access
```

See the *Dell PowerConnect ArubaOS 5.0 Command Line Reference Guide* for detailed information on the command options.

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. By default, these yet-unauthorized APs are put into the temporary AP group **authorization-group** 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 the remote AP 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.

Add or Edit 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** to save your changes.

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>
```

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 controller and takes some action based on that identification. You apply these ACLs to user roles or uplink ports.



Note: To configure firewall policies, you must install the PEFNG license.

For more information about ACLs and firewall policies, see [“Configuring the fallback mode” on page 169](#).

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 30 Sample Split Tunnel Environment

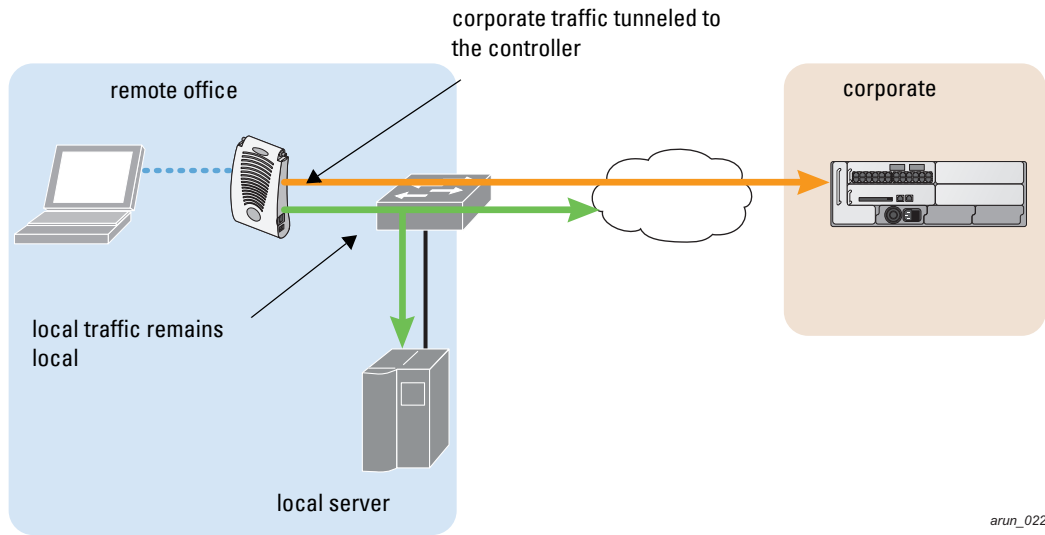


Figure 30 displays corporate traffic is GRE tunneled to the controller through a trusted tunnel and local traffic is source NATed and bridged on the wired interface based on the configured user role and session ACL.

Configuring Split Tunneling

To configure split tunneling:

- Define a session ACL that forwards only corporate traffic to the controller.
 - Configure a netdestination for the corporate subnets.
 - Create rules to permit DHCP and corporate traffic to the corporate controller. When specifying the action that you want the controller to perform on a packet that matches the specified criteria, “permit” implies tunneling, which is used for corporate traffic, and “route” implies local bridging, which is used for local traffic.

You must install the PEFNG license in the controller. For information about user roles and policies, see [Chapter 10, “Roles and Policies”](#).

- Apply the session ACL to a user role.
- Configure the AAA profile.

The AAA profile defines the authentication method and the default user role for authenticated users. The configured user role contains the split ACL.



Note: 802.1x and PSK authentication is supported when configuring split tunnel mode.

- Configure the virtual AP profile:

When configuring the virtual AP profile, you specify which AP group or AP the profile applies to.

 - Set the VLAN used for split tunneling. Only one VLAN can be configured for split tunneling; VLAN pooling is not allowed.
 - When specifying the use of a split tunnel configuration, use “split-tunnel” forward mode.
 - Create and apply the applicable SSID profile.



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 [“Configuring Profiles” on page 105 in Chapter 6, “Remote Access Points”](#).

- Optionally, create a list of network names resolved by 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.

Configuring the Session ACL

First you need to configure the session ACL. By applying this policy, local traffic remains local, and corporate traffic is forwarded (tunneled) to the controller.

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 **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 **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.
7. 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.
8. Under **Destination**, select the alias you just created.
9. Under **Service**, select **any**.
10. Under **Action**, select **permit**.
11. Click **Add**.
12. 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**.
13. Click **Apply**.
14. 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**.
15. Click **Apply**.

Using the CLI

```

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 ACL for restricted LD homepage access

A user in split or bridge role using a remote AP (RAP) can log on to the local debug (LD) homepage (for example, `http://10.10.10.10`) 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 new `localip` keyword in the in the user role ACL.



Note: You will require the PEF license to use this feature. See [Chapter 27, “Software Licenses” on page 527](#) 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.

Using CLI

Use the `localip` keyword in the user role ACL.

By default, all users have an ACL entry of type `any any deny`. 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.

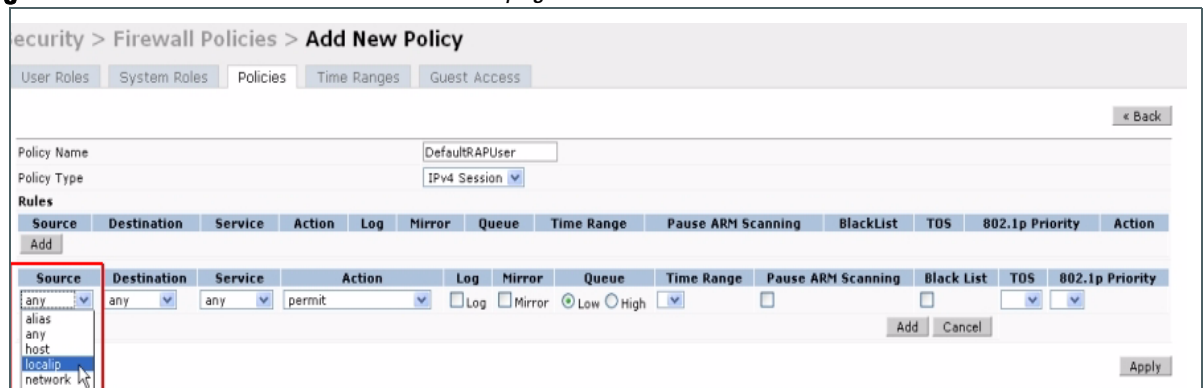
Example:

```
ip access-list session logon-control
  user localip svc-http deny
  user any permit
```

Using 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 31 Enable Restricted Access to LD Homepage



Configuring the AAA Profile and the Virtual AP Profile

After you configure the session ACL, you define the AAA profile and virtual AP used for split tunneling. When defining the AAA parameters, specify the previously configured user role that contains the session ACL used for split tunneling.

Using 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**.

If you need to create an authentication server group, select **new** and enter the appropriate parameters.

Using CLI

```
aaa profile <name>
  authentication-dot1x <dot1x-profile>
  dot1x-default-role <role>
  dot1x-server-group <group>
```

Configuring split tunneling in the virtual AP profile

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**.



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, 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**.

Using the CLI to configure split tunneling in the virtual AP profile

```
wlan ssid-profile <profile>
  essid <name>
  opmode <method>

wlan virtual-ap <profile>
  ssid-profile <name>
  forward-mode split-tunnel
  vlan <vlan id>
  aaa-profile <profile>
```

```
ap-group <name>
  virtual-ap <profile>
```

or

```
ap-name <name>
virtual-ap <profile>
```

Using the WebUI to list the corporate DNS servers

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**.

Using the CLI to list the corporate DNS servers

```
ap system-profile <profile>
  dns-domain <domain name>
```

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 [Chapter 29, “Voice and Video”](#).

Uplink Bandwidth Reservation

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.

The following must be noted 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 will be treated equal.

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 also on 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 CLI or the WebUI.

Using the WebUI

To configure bandwidth reservation

1. Navigate to **Configuration > Advanced Services > All Profiles**.
2. Under **Profiles**, navigate to **AP > AP System Profile**. Create a new AP system profile to configure bandwidth reservation or edit an existing AP system profile.
3. Under the **Profile Details** page, specify bandwidth reservation values.

Figure 32 Uplink Bandwidth Reservation

Remote-AP uplink total bandwidth	<input type="text" value="1024"/> kbps	RAP bw reservation 1	aclname <input type="text" value="voice"/> bwvalue <input type="text" value="128"/> prio <input type="text" value="1"/>
RAP bw reservation 2	aclname <input type="text"/> bwvalue <input type="text"/> prio <input type="text"/>	RAP bw reservation 3	aclname <input type="text"/> bwvalue <input type="text"/> prio <input type="text"/>
Heartbeat DSCP	<input type="text" value="0"/>	Session ACL	<input type="text" value="ap-uplink-acl"/> <input type="button" value="v"/>

Using CLI

```
(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   1524      370962    0         1524      0         0         0         0
```


The Dell 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. This chapter includes the following topics:

- [“Mesh Access Points” on page 189](#)
- [“Mesh Links” on page 191](#)
- [“Mesh Profiles” on page 192](#)
- [“Mesh Solutions” on page 195](#)
- [“Before You Begin” on page 198](#)
- [“Mesh Radio Profiles” on page 199](#)
- [“RF Management \(802.11a and 802.11g\) Profiles” on page 204](#)
- [“Mesh High-Throughput SSID Profiles” on page 211](#)
- [“Mesh Cluster Profiles” on page 214](#)
- [“Ethernet Ports for Mesh” on page 219](#)
- [“Provisioning Mesh Nodes” on page 222](#)
- [“AP Boot Sequence” on page 224](#)
- [“Verifying the Network” on page 225](#)
- [“Remote Mesh Portals” on page 225](#)

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.



Note: Remote Mesh Portal only (not Mesh points) is supported on RAP-5WN. Mesh is not supported on RAP-2WG.

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, that mesh node will not deliver WLAN services to its clients.

For mesh as well as 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.

By default, APs operate as thin APs, which means 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 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 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” on page 219 or “Provisioning Mesh Nodes” on page 222.

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) 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” on page 193](#).

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 33](#) shows two mesh clusters and their relationship to the controller.

Figure 33 *Sample Mesh Clusters*



Mesh Links

In simple terms, 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, specifically 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: Through out the rest of this chapter, the term “uplink” is also 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 unavailable, 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 re-establish 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 38](#) 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. The mesh portal advertises a cost of 0, while all other mesh nodes advertise a cumulative cost based on the parent mesh node.

Table 38 *Mesh Link Metric Computation*

Metric	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 quality	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.
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 will use a legacy path if no high-throughput path is available.

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, Dell recommends the default values for most deployments. For information, see [“Metric algorithm” on page 200](#).

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 Profile

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 will 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.

Dell provides a “default” version of the mesh cluster profile. You can use the “default” version or create a new instance of a profile which you can then edit as you need. You can configure a maximum of 16 mesh cluster profiles on a mesh node. For details about configuring mesh cluster profiles, see [“Mesh Cluster Profiles” on page 214](#).

Mesh Radio Profile

Dell provides a “default” version of the mesh radio profile. You can use the “default” version or create a new instance of a profile which you can then edit as you need. The mesh radio profile allows you to specify the set of rates used to transmit data on the mesh link. For information about configuring mesh radio profiles, see [“Mesh Radio Profiles” on page 199](#).

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 [“RF Management \(802.11a and 802.11g\) Profiles” on page 204](#).

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” on page 214](#)).

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 will 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 will 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 will be overridden by any channel changes from the mesh portal.

A mesh point will take 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**.

The ARM power adjustment feature does not apply to all ARM-enabled Mesh portals. Indoor mesh portals can take advantage of this feature to adjust power settings according to their ARM profiles, but outdoor mesh portals will continue to run at their configured power level to maximize their range. It is also important to note that 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 will tune to this channel, form the link, and will 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, and once the mesh network has formed and all the mesh points have selected a portal (or upstream mesh point), those mesh points will 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 will 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 208](#).

High-Throughput 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 207](#).

Mesh High-Throughput SSID Profile

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” on page 211](#).

Wired AP Profile

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 [“Ethernet Ports for Mesh” on page 219](#)

Mesh Recovery Profile

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.

Note: 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 controllers' configuration file and is unique to that master controller. If necessary, you can transfer your configuration to another controller. If you do this, make sure your new mesh cluster is running and you have re-provisioned the mesh nodes before deleting your previous configuration. The APs will 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 will not work on the mesh node until you re-provision it with the new master controller.

Mesh 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 34](#) shows the wireless backhaul between the mesh portal to the mesh point that services the wireless clients.


Figure 34 *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 35](#) shows a single-hop point-to-point deployment.


Figure 35 *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 36](#) shows a single-hop point-to-multipoint deployment.


Figure 36 *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 37](#) 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 37 *Sample High-Availability Deployment*



Before You Begin

Dell recommends the following when planning and deploying a mesh solution:

Pre-Deployment Considerations

- Ensure the controller has Layer-2/3 network connectivity to the network segment where the mesh portal will be installed.
- 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.
- Stage the APs before deployment. Identify the location of the APs, configure them for mesh, and provision and verify connectivity them before deploying them in a live network.
- 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, Dell recommends using 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 1024 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.

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.



Note: 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. Dell recommends creating 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 affect. 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.

AP-12x Specific Considerations

The W-API20 SeriesI 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, the same wired AP profile will control both enet0 and enet1.

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 affect immediately. You do not reboot the controller or the AP.

Managing Mesh Profiles via the WebUI

Use the following procedures to define and manage mesh radio profiles via the WebUI.

Creating a New 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 the **Edit** button by the AP group name for which you want to configure the new mesh radio profile.
 - If you selected the **AP Specific** tab, click the **Edit** button 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. 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. You cannot use spaces in radio profile names.
4. Configure your desired mesh radio settings. [Table 39](#) describes the parameters you can configure in the mesh radio profile

Table 39 Mesh Radio Profile Configuration Parameters

Parameter	Description
Mesh radio profile	Select an existing radio profile to modify or create a new radio profile. The radio profile can have a maximum of 32 characters. Default: Mesh radio profile named "default."
Maximum Children	Indicates the maximum number of children a mesh node can accept. Default: 64 children. The range is 1–64.
Maximum Hop Count	Indicates the maximum hop count from the mesh portal. Default: 8 hops. The range is 1–32.
Heartbeat threshold	Indicates the maximum number of heartbeat messages that can be lost between neighboring mesh nodes. Default: 10 missed heartbeats. The range is 1–255.
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.
Metric algorithm	Use this setting to optimize operation of the link metric algorithm. Specifies the algorithm used by a mesh node to select its parent. Available options are: <ul style="list-style-type: none"> 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. <p>NOTE: Dell recommends using the default value.</p> Default: distributed-tree-rssi.

Table 39 Mesh Radio Profile Configuration Parameters (Continued)

Parameter	Description
Reselection mode	<p>Use this setting to optimize 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.</p> <p>Available options are:</p> <ul style="list-style-type: none"> 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). Dell recommends using this default startup-subthreshold value. subthreshold-only—Connected mesh nodes evaluate alternative links only if the existing uplink becomes a sub-threshold link. <p>NOTE: 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 will seek 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 will continue 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.</p>
Retry Limit	<p>Indicates the number of times a mesh node can re-send a packet.</p> <p>Default: 4 times. The range is 0– 15.</p>
RTS Threshold	<p>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.</p> <p>Default: 2,333 bytes. The range is 256– 2,346.</p>
802.11a Transmit Rates	<p>Indicates the transmit rates for the 802.11a radio.</p> <p>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.</p> <p>To modify transmit rates, do one of the following:</p> <ul style="list-style-type: none"> 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. <p>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.</p>
802.11g Transmit Rates	<p>Indicates the transmit rates for the 802.11g radio.</p> <p>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.</p> <p>To modify transmit rates, do one of the following:</p> <ul style="list-style-type: none"> 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. <p>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.</p>
Mesh Private VLAN	<p>A VLAN ID for control traffic between an remote mesh portal and mesh nodes. This VLAN ID must not be used for user traffic.</p> <p>Range: 0–4094. Default: 0 (disabled).</p> <p>For further information on configuring a remote mesh portal, see “Remote Mesh Portals” on page 225</p>

Table 39 Mesh Radio Profile Configuration Parameters (Continued)

Parameter	Description
Allowed VLANs on Mesh Link	List the VLAN ID numbers of VLANs allowed on the mesh link.
BC/MC Rate Optimization	<p>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.</p> <p>When the Multicast Rate Optimization feature is enabled, 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.</p> <p>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 will be 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.</p> <p>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.</p> <p>Default: Enabled.</p>

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 Profile to a Mesh AP or AP Group

- 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.
- Under the Profiles list, expand the **Mesh** menu, then select **Mesh radio profile**.
- In the **Profile Details** window pane, click the **Mesh radio profile** drop-down list and select the desired mesh radio profile from the list.
- 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.

Editing a Profile

- 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.
- In the Profiles list, expand the **Mesh** menu, then select **Mesh radio 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.
- Change the mesh radio settings as desired. [Table 39](#) describes the parameters you can configure in the mesh radio profile.
- Click **Apply** to save your changes.

Deleting a Profile

Use the following procedure to delete an existing mesh radio profile via the WebUI. You can delete a mesh radio profile only if no other APs or AP groups are using that profile.

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 the **Delete** button by the name of the profile you want to delete.

Managing Mesh Profiles using 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 Profile

Configuration details and any default values for each of these parameters are described in [Table 39 on page 200](#). 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.

```
ap mesh-radio-profile <profile-name>
  a-tx-rates
  allowed-vlans
  children <children>
  clone <source-profile-name>
  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
  metric-algorithm {best-link-rssi|distributed-tree-rssi}
  mpv <vlan-id>
  no
  rts-threshold <rts-threshold>
  tx-power <tx-power>>
```

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.

```
ap mesh-radio-profile <profile-name>
  clone <source-profile-name>
```

Viewing Profile Settings

To view a complete list of mesh radio profiles and their status:

```
show ap mesh-radio-profile
```

To view the settings of a specific mesh radio profile:

```
show ap mesh-radio-profile <name>
```

Assigning a Profile to an 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.

```
ap-group <group>
  mesh-radio-profile <profile-name> priority <priority>
```

To associate a mesh radio profile with an individual AP:

```
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.

group2-cluster1 has a priority of 10, and cluster2 has a priority of 5.

```
ap-group group1
  mesh-cluster-profile cluster1 priority 5
  mesh-cluster-profile cluster2 priority 10
```

```
ap-group2
  mesh-cluster-profile cluster1 priority 10
  mesh-cluster-profile cluster2 priority 5
```

Deleting a Mesh Radio Profile

If no AP or AP group is using a mesh radio profile, you can delete that profile using the **no** parameter:

```
no ap mesh-radio-profile <profile-name>
```

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. 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.

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.

Managing 802.11a/802.11g Profiles Via the WebUI

Use the following procedures to define and manage 802.11a and 802.11g RF management profiles via the WebUI.

Creating 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 a new RF management profile.
 - If you selected **AP Specific**, click the **Edit** button by the AP for which you want to create a new 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. If you selected **802.11a radio profile**, click the **802.11a radio profile** drop-down list in the **Profile Details** window pane and select **NEW**.
-or-
If you selected **802.11g radio profile**, click the **802.11g radio profile** drop-down list in the **Profile Details** window pane and select **NEW**.
4. Enter a name for your new 802.11a or 802.11g radio profile.
5. Configure the radio settings described in [Table 40](#), then click **Apply** to save your settings. The profile name appears in the Profile list with your configured settings.

Table 40 802.11a/802.11g RF Management Configuration Parameters

Parameter	Description
ARM profile	Adaptive Radio Management (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 208.
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 207.
Radio Enable	Enable transmissions on this radio band.
Mode	Access Point operating mode. Available options are: <ul style="list-style-type: none"> ● am-mode: Air Monitor mode ● ap-mode: Access Point mode ● apm-mode: Access Point Monitor mode ● sensor-mode: RFprotect sensor 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.
Channel	Transmit channel for this radio.
Beacon Period	Beacon Period for the AP in msec. The minimum value is 60 msec, and the default value is 100 msec.
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.
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.
Spectrum Load Balancing Domain	Enter a spectrum load balancing domain name to manually create RF neighborhoods. Use this option to create RF neighborhood information for networks that have disabled Adaptive Radio Management (ARM) scanning and channel assignment. <ul style="list-style-type: none"> ● If spectrum load balancing is enabled in a 802.11g radio profile but the spectrum load balancing domain is <i>not</i> defined, ArubaOS uses the ARM feature to calculate RF neighborhoods. ● If spectrum load balancing is enabled in a 802.11g radio profile and a spectrum load balancing domain <i>is also</i> defined, AP radios belonging to the same spectrum load balancing domain will be considered part of the same RF neighborhood for load balancing, and will not recognize RF neighborhoods defined by the ARM feature.

Table 40 802.11a/802.11g RF Management Configuration Parameters (Continued)

Parameter	Description
Spectrum Load Balancing	<p>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.</p> <p>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 149.</p>
RX Sensitivity Tuning Based Channel Reuse	<p>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.</p> <p>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 "RX Sensitivity Tuning Based Channel Reuse" on page 149.</p> <p>NOTE: Do not enable the Channel Reuse feature if Non 802.11 Interference Immunity is set to level 3 or higher. A level-3 to level-4 Noise Immunity setting is not compatible with the Channel Reuse feature.</p>
RX Sensitivity Threshold	<p>RX sensitivity tuning based channel reuse threshold, in - dBm.</p> <p>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.</p> <p>If the value for this parameter is set to zero, the feature will automatically determine an appropriate threshold.</p>
Non 802.11 Interference Immunity	<p>(for 802.11g profiles only)</p> <p>Set a value for 802.11 Interference Immunity. This parameter sets the interference immunity on the 2.4 GHz band.</p> <p>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.</p> <p>The levels for this parameter are:</p> <ul style="list-style-type: none"> ● 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. <p>NOTE: Do not raise the noise immunity feature's default setting if the RX Sensitivity Tuning Based Channel Reuse feature is also enabled. A level-3 to level-5 Noise Immunity setting is not compatible with the Channel Reuse feature.</p>
Enable CSA	<p>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.</p>
CSA Count	<p>Number of channel switch announcements that must be sent prior to switching to a new channel. The default CSA count is 4 announcements.</p>
Management Frame Throttle Interval	<p>Averaging interval for rate limiting mgmt frames from this radio, in seconds. A management frame throttle interval of 0 seconds disables rate limiting.</p>
Management Frame Throttle Limit	<p>Maximum number of management frames that can come in from this radio in each throttle interval.</p>

Table 40 802.11a/802.11g RF Management Configuration Parameters (Continued)

Parameter	Description
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 override functions are always enabled, regardless of whether or not this check box is selected.
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.
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 24–58km, 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.

Assigning an 802.11a/802.11g Profile

Use the following procedure to assign an 802.11a or 802.11g RF management profile to an AP group or individual AP via the WebUI.

- 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
- Under the Profiles list, expand the **RF management** menu.
- To select a **802.11a radio profile** for an AP or AP group, click **802.11a radio profile**. In the **Profile Details** window pane, click the **802.11a radio profile** drop-down list and select the desired profile from the list
-or-
To select a **802.11g radio profile** for an AP or AP group, click **802.11g radio profile**. In the **Profile Details** window pane, click the **802.11g radio profile** drop-down list and select the desired profile from the list
- 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.

- 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.
3. To reference a new high-throughput profile for an *802.11a* RF management profile, expand the **802.11a radio profile** menu, then select **High-throughput radio profile.**
-or-
To reference a new high-throughput profile for an *802.11g* RF management profile, expand the **802.11g radio profile** menu, then select **High-throughput radio profile.**
4. 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 **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.
 - d. Click **Apply** to save your settings.
5. The high-throughput profile name 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 to 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.

Editing an 802.11a/802.11g 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 name using the 802.11a or 802.11g RF management profile you want to edit.
 - If you selected **AP Specific**, click the **Edit** button by the AP using the 802.11a or 802.11g RF management profile you want to edit.
2. Under the Profiles list, expand the **RF** menu.
3. To edit an **802.11a radio profile** for an AP or AP group, click **802.11a radio profile**. In the **Profile Details** window pane, click the **802.11a radio profile** drop-down list and select the desired profile from the list
-or-
To select a **802.11g radio profile** for an AP or AP group, click **802.11g radio profile**. In the **Profile Details** window pane, click the **802.11g radio profile** drop-down list and select the desired profile from the list
4. Change the profile settings as desired. [Table 40](#) describes the parameters you can configure in the mesh 802.11a or 802.11g RF management profile.
5. Click **Apply** to save your changes.

Deleting a Profile

You can delete a mesh high-throughput SSID profile only if no APs or AP groups are associated with that profile. To delete a 802.11a or 802.11g radio profile via the WebUI.

1. Navigate to the **Configuration > Advanced Services > All Profiles** window.
2. Expand the **RF** 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 40 on page 205](#). 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>
  arm-profile
  beacon-period
  channel
  channel-reuse
  channel-reuse-threshold
  clone
  csa
  csa-count
  disable-arm-wids-function
  dot11b-protection (for 802.11b profiles only)
  dot11h
  high-throughput-enable
  ht-radio-profile
  interference-immunity (for 802.11g profiles only)
  maximum-distance <maximum-distance>
  mgmt-frame-throttle-interval
  mgmt-frame-throttle-limit
  mode
  no
  radio-enable
  spectrum-load-balancing
  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>
```

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 Profiles via the WebUI

Use the following procedures to manage your high-throughput SSID profiles via 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 the **Edit** button by the AP group for which you want to create the new high-throughput SSID profile.
 - If you selected **AP Specific**, click **Edit** button 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 high-throughput SSID described in [Table 41](#), then click **Apply** to save your settings. The profile name appears in the Mesh High-throughput SSID Profile list with your configured settings.

Table 41 Mesh High-Throughput SSID Profile Configuration Parameters

Parameter	Description
Mesh high-throughput SSID profile	Enter the name of an existing mesh high-throughput SSID profile to modify that profile, or enter a new name or create a new mesh high-throughput profile. The mesh high-throughput profile can have a maximum of 32 characters. To view existing high-throughput SSID radio profiles, use the command: <code>show ap mesh-radio-profile</code> .
High throughput enable (SSID)	Enable or disable high-throughput (802.11n) features on this SSID. This parameter is enabled by default.
MPDU Aggregation	Enable or disable MAC protocol data unit (MPDU) aggregation. High-throughput mesh 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 transmitted A-MPDU size	Maximum size of a transmitted aggregate MPDU, in bytes. Range: 1576–65535
Max received A-MPDU size	Maximum size of a received aggregate MPDU, in bytes. Allowed values: 8191, 16383, 32767, 65535.

Table 41 Mesh High-Throughput SSID Profile Configuration Parameters (Continued)

Parameter	Description
Min MPDU start spacing	Minimum time between the start of adjacent MPDUs within an aggregate MPDU, in microseconds. Allowed values: 0 (No restriction on MPDU start spacing), .25 μ sec, .5 μ sec, 1 μ sec, 2 μ sec, 4 μ sec.
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 1–15; 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: 0–15.
Legacy stations	Allow or disallow associations from legacy (non-HT) stations. By default, this parameter is enabled (legacy stations are allowed).
40 MHz channel usage	Enable or disable the use of 40 MHz channels. This parameter is enabled by default.
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.

Assigning a Profile to an AP Group

- 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 SSID profile.
 - If you selected **AP Specific**, click the **Edit** button by the AP to which you want to assign a new high-throughput SSID profile
- Under the Profiles list, expand the **Mesh** menu, then select **Mesh High-throughput SSID profile**.
- In the **Profile Details** window pane, click the **Mesh High-throughput SSID profile** drop-down list and select the desired profile from the list.
- 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

- 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.
- 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 41](#) describes the parameters you can configure in this profile.
5. Click **Apply** to save your changes.

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 the **Delete** button by the name of the profile you want to delete.

Managing high-throughput SSID profiles using the CLI

You must be in config mode to create, modify or delete a mesh 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 41 on page 211](#). 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

```
ap mesh-ht-ssid-profile <profile-name>
  40MHz-enable
  clone
  high-throughput-enable
  legacy-stations
  max-rx-a-mpdu-size
  max-tx-a-mpdu-size
  min-mpdu-start-spacing
  mpdu-agg
  no
  short-guard-intvl-40Mhz
  supported-mcs-set
```

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:

```
ap-group <group> mesh-ht-ssid-profile <profile-name>
```

To associate a mesh radio profile with an individual AP:

```
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:

```
show ap mesh-ht-ssid-profile
```

To view the settings of a specific high-throughput profile:

```
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:

```
no ap mesh-ht-ssid-profile <profile-name>
```

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” profile, are not applied until you provision your APs for mesh.

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. Dell recommends creating 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” on page 222](#) for more information.

Deployments with Multiple Mesh Cluster Profiles

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 will use 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 a sample configuration, see “show ap mesh topology” on page 225.

Managing Mesh Cluster Profiles via the WebUI

Use the following procedures to define and manage mesh cluster profiles via 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 the **Edit** button by the AP group name for which you want to create the new mesh cluster profile.
 - If you selected **AP Specific**, click the **Edit** button 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 42](#), then click **Apply** to save your settings.

Table 42 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, <i>do not use spaces in the mesh cluster name</i> , as this may cause errors in mesh points associated with that mesh cluster. To view existing mesh cluster profiles, use the CLI command: <code>show ap mesh-cluster-profile</code> . 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, <i>each mesh cluster profile must use the same band</i> .
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). Dell recommends selecting 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.

Table 42 Mesh Cluster Profile Configuration Parameters (Continued)

Parameter	Description
Priority	<p>Indicates the priority of the cluster profile.</p> <p>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</p> <p>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).</p> <p>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.</p> <p>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.</p>
Cluster Name	<p>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.</p> <p>NOTE: Each mesh cluster profile should have a unique MSSID. Configure a new MSSID before you apply the mesh cluster profile.</p> <p>To view existing mesh cluster profiles, use the command: <code>show ap mesh-cluster-profile</code>.</p> <p>A mesh portal chooses the best cluster profile and provisions it for use. A mesh point can have a maximum of 16 cluster profiles.</p> <p>Default: Mesh cluster named “Dell-mesh.”</p>
RF Band	Indicates the band for mesh operation for multiband radios. Select a or g .

Associating a 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 via 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 the **Edit** button by the AP group name to which you want to assign a new mesh cluster profile.
 - If you selected **AP Specific**, click the **Edit** button 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 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” on page 222](#).

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 41](#) describes the parameters you can configure in the mesh high-throughput SSID profile.

 **Note:** 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 Using 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 42 on page 215](#). 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.

```
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`.

```
ap mesh-cluster-profile cluster1
  cluster corporate
  opmode wpa2-psk-aes
  wpa-passphrase mesh_123
  rf-band a
```

```
ap mesh-cluster-profile cluster2
  cluster corporate
  opmode wpa2-psk-aes
  wpa-passphrase mesh_123
  rf-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.

```
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:

```
show mesh-cluster-profile
```

To view the settings of a specific mesh cluster profile:

```
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:

```
ap-group <group> mesh-cluster-profile <profile-name>
```

To associate a mesh cluster profile to an individual AP in a single-cluster deployment:

```
ap-name <name> mesh-cluster-profile <profile-name>
```

To associate a mesh cluster profile to an AP group in a multiple-cluster deployment:

```
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

```
ap-name <name>
  mesh-cluster-profile <profile-name> priority <priority>
```

Example:

```
ap-group group1
  mesh-cluster-profile cluster1 priority 5
  mesh-cluster-profile cluster2 priority 10
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:

```
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:

```
no ap mesh-cluster-profile <profile-name>
```

Ethernet Ports for Mesh

If you are using 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.

Note: 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 the AP-70 or AP-12x, note the following requirements:

- If the AP is 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 the AP is configured as a mesh point, the same wired AP profile will control both enet0 and enet1.

Configure 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.

```
ap wired-ap-profile <profile>
    forward-mode bridge
    wired-ap-enable
```

Optionally, you can configure the following wired AP profile settings:

```
ap wired-ap-profile <profile>
    switchport mode {access | trunk}
    switchport access vlan <vlan>
    switchport trunk native vlan <vlan>
    switchport trunk allowed vlan <vlan>
    trusted
```

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 AP-70 or AP-12x.

You configure secure jack operation in the wired AP profile.



Note: 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.

Configuring Secure Jack via the WebUI

Use the following procedure to configure secure jack operation 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 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** to save your settings.

Configuring Secure Jack via the CLI

To configure secure jack operation via the command-line interface, access the CLI in config mode and issue the following commands:


```
ap wired-ap-profile <profile>
    forward-mode tunnel
    wired-ap-enable
```

Optionally, you can configure the following wired AP profile settings:

```
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:

 **Note:** Dell recommends 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, Dell recommends a value of 10,000.
- **Bootstrap threshold**—Number of consecutive missed heartbeats (heartbeats are sent once per second) before the AP reboots. The default is 9 missed heartbeats. If you must modify this setting, Dell recommends a value of 5,000.

When the controller comes back online, the affected mesh nodes (mesh portals and mesh points) will rebootstrap; however, the mesh link is not affected and will continue to be up.

Modifying the AP System Profile via 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**.

Modifying the AP System Profile via the CLI

To modify the AP system profile via the command-line interface, access the CLI in config mode and issue the following commands:

```
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” on page 199](#) 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” on page 190](#). 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 will be 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 will use the recovery profile to establish a link to the master controller; data forwarding does not take place.

Note: If you create a new mesh cluster profile for an existing deployment, you must re-provision the AP for the new profile to take affect. 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.

This section describes the following topics:

- [“Outdoor AP Parameters” on page 222](#)
- [“Provisioning Caveats” on page 223](#)
- [“Provisioning Mesh Nodes via the WebUI” on page 223](#)

Outdoor AP Parameters

If you are using outdoor APs and planning an outdoor mesh deployment, you can enter the following outdoor parameters when provisioning the AP:

- Latitude and longitude coordinates of the AP. These location identifiers allow you to more easily locate the AP for inventory and troubleshooting purposes.
- Altitude, in meters, of the AP.
- Antenna bearing to determine horizontal coverage.
- Antenna angle for optimum antenna coverage.

Note: The above parameters apply to all outdoor APs, not just outdoor APs configured for mesh.

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” on page 225.
- 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 via the WebUI

Reprovisioning the AP causes it to automatically reboot. The easiest way to provision a mesh node is to use the **Provisioning** window in the WebUI. The following procedure describes the process to provision a mesh portal or mesh node via the WebUI or CLI. To provision a remote mesh portal, see “[Remote Mesh Portals](#)” on page 225.

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)
5. Click **Apply and Reboot**. After the controller reboots, mesh cluster profiles are extracted from the AP group and the AP name.

Provisioning Mesh Nodes via the CLI

When you use the command-line interface to reprovision a mesh node, you may also provision other AP settings. To provision a remote mesh portal, see [“Remote Mesh Portals” on page 225](#).

Access the CLI in config mode and issue the following commands:

```
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:

```
provision-ap
  read-bootinfo ap-name <name>
  mesh-role {mesh-point|mesh-portal|remote-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>
```

AP Boot Sequence

The information in this section describes the boot sequence for mesh APs. Depending on their configured role, the AP performs a slightly different boot sequence.

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.

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 then uses Dell Discovery Protocol (ADP) to discover the master controller. 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.



Note: 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 will not trigger “wireless-bridging” events for packets transmitted between mesh nodes.

Verifying the Network

After provisioning the mesh APs, ensure that the mesh network is up and operating correctly.

To view your network via the WebUI, navigate to the one of the following windows:

- **Monitoring > Network > All Mesh Nodes**
- **Monitoring > Network > controller> Mesh Nodes**

To view your network via the command line interface, use the following commands:

- `show ap mesh active`
- `show ap mesh topology`

Remote Mesh Portals

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 is used to secure 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 [Chapter 27, “Software Licenses” on page 527.](#)

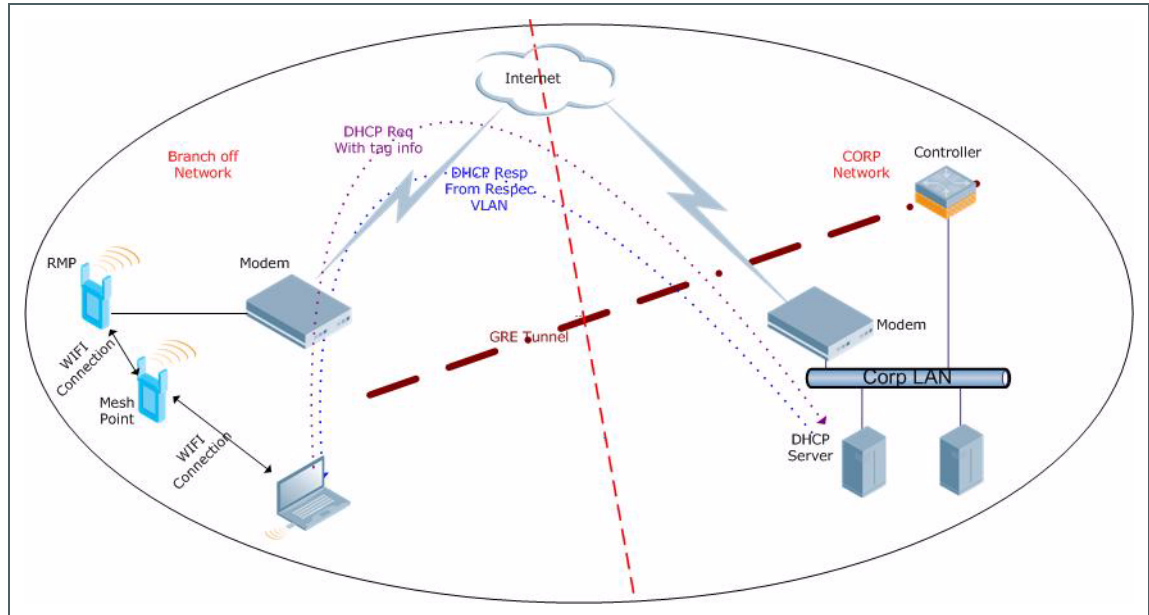
How RMP Works

When a client at the branch office associates to a split VAP, 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 set up for the RMP on the controller, the VLAN of the tunnel should be the MPV. A DHCP pool for the MPV should be set up on the controller. 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.

Figure 38 Working of RMP



Creating a Remote Mesh Portal via 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 154](#).

Wired ports on remote mesh portals can be configured in either bridge or split-tunnel forwarding mode. There are, however, 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.

Provisioning 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, click the **Mesh Role** drop-down list and select **Remote Mesh Portal**.

Figure 39 Provisioning an AP as a Remote Mesh Portal

The screenshot shows the configuration page for provisioning an AP as a Remote Mesh Portal. The 'Antenna Selection' section has 'Internal/Included Antenna' selected. Under 'Authentication Method', 'Remote AP' is set to 'Yes'. The 'Remote AP Authentication Method' is 'Pre-shared Key', with an IKE PSK field and a 'Confirm IKE PSK' field. In 'User credential assignment', 'Use Automatic Generation' is checked, and 'Global User Name/Password' is selected. The 'User Name' is 'aruba123' and the 'Password' is masked. There are 'Generate' buttons for both. The 'PPPoE Parameters' section is collapsed. 'Master Discovery' has 'Host Controller IP Address' set to '10.4.46.234' and 'Master Controller IP Address/DNS name' also set to '10.4.46.234'. 'IP Settings' has 'Obtain IP Address Using DHCP' selected. 'FQLN Mapper' has 'Remove FQLN' unchecked and 'Main Campus' selected for 'Campus'. The 'AP List' table at the bottom shows one entry with IP '1.1.1.1', AP Name 'ap65_1_22', and AP Group 'pixel8'. The 'Mesh Role' dropdown is open, showing options: 'none', 'Radio', 'Mesh Point', 'Mesh Portal', and 'Remote Mesh Portal' (highlighted in red). 'Apply and Reboot' and 'Cancel' buttons are at the bottom right.

Defining the Mesh Private VLAN

Edit the mesh radio profile for the remote mesh portal and 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 will come up in that Mesh Private Vlan. This mesh private VLAN must not be used as a VLAN for any other virtual AP.

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 remote mesh portal AP group with the profile you want to edit.
 - If you selected the **AP Specific** tab, click the **Edit** button by the remote mesh portal with the profile you want to edit.
2. In 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 name of the profile you want to edit.
4. Set the **Mesh Private VLAN** parameter to define a VLAN ID (0–4094) for control traffic between an remote mesh point and mesh nodes.
5. 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.

Selecting a Mesh Radio Profile

Use the following procedure to select a mesh radio profile for a remote 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.

Selecting an RF Management Profile

Use the following procedure to select an RF management profile for a remote 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 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.
3. To select a **802.11a radio profile** for an AP or AP group, click **802.11a radio profile**. In the **Profile Details** window pane, click the **802.11a radio profile** drop-down list and select the desired profile from the list
-or-
To select a **802.11g radio profile** for an AP or AP group, click **802.11g radio profile**. In the **Profile Details** window pane, click the **802.11g radio profile** drop-down list 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.

Adding a Mesh Cluster Profile

Use the following procedure to add a mesh cluster profile to a remote 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 name to which you want to assign a new mesh cluster profile.
 - If you selected **AP Specific**, click the **Edit** button 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.

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.



Note: 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.

5. Click **Add** to add the mesh cluster profile to the AP group.

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 will 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” on page 171.

Configuring the VLAN ID of the Virtual AP Profile

The VLAN of this Virtual AP must have the same VLAN ID as the mesh private VLAN.

1. Navigate to **Configuration > Wireless > AP Configuration** window. Select either the **AP Group** or **AP Specific** tab. Click the **Edit** button by the applicable AP group name or AP name with the virtual AP profile you want to configure.
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 “Dell-ap” 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** window, click 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 **Profile Details** window pane to display the configuration parameters for this profile.
 6. In the **Profile Details** window:
 - a. Make sure **Virtual AP enable** is selected.
 - b. From the **VLAN** drop-down menu, select the VLAN ID for the mesh private VLAN.
 - c. From the **Forward mode** drop-down menu, select **split-tunnel**.
 - d. Click **Apply**.

Provisioning a Remote Mesh Portal via 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.

```
provision-ap
  read-bootinfo ap-name <name>
  mesh-role remote-mesh-portal
  reprovision ap-name <name>
```

Additional Information

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 case, the portal performs the default operation that is to bridge 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.

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.

Important Points to Remember

- In order 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
<http://technet2.microsoft.com/windowsserver/en/technologies/ias.mspx>
Active Directory
<http://technet2.microsoft.com/windowsserver/en/technologies/featured/ad/default.mspx>

This chapter describes the following topics:

- “Servers and Server Groups” on page 231
- “Configuring Servers” on page 232
- “Internal Database” on page 236
- “Server Groups” on page 239

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)

Additionally, you can use the controller's internal database to authenticate users. You create entries in the database for users and their passwords and 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 40](#) 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 40 *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 are using 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 Servers

This section describes how to configure RADIUS, LDAP, TACACS+ and Windows external authentication servers and the internal database on the controller.

Configuring a RADIUS Server

[Table 43](#) describes the parameters you configure for a RADIUS server.

Table 43 *RADIUS Server Configuration Parameters*

Parameter	Description
Host	IP address of the authentication server. 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

Table 43 RADIUS Server Configuration Parameters (Continued)

Parameter	Description
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 <code>ip radius nas-ip ipaddr</code> command. Default: N/A
Use MD5	Use MD5 hash of cleartext password. Default: disabled
Mode	Enables or disables the server. Default: enabled

In 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 43](#). Select the **Mode** checkbox to activate the authentication server.
5. Click **Apply** to apply the configuration.



Note: The configuration does not take effect until you perform this step.

In the CLI

```
aaa authentication-server radius <name>
  host <ipaddr>
  key <key>
  enable
```

RADIUS Server Authentication Codes

A configured RADIUS server will return the following standard response codes.

Table 44 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.)

Configuring an LDAP Server

Table 45 describes the parameters you configure for an LDAP server.

Table 45 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 need not have write privileges but the user should 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 which contains the entire user database to use. Default: N/A
Filter	Filter that should be applied to search of the user in the LDAP database (default filter string is: <code>!(objectclass=*)</code>). Default: N/A
Key Attribute	Attribute that should be used as a key in search for the LDAP server. For Active Directory, the value is <code>sAMAccountName</code> . Default: <code>sAMAccountName</code>
Timeout	Timeout period of a LDAP request, in seconds. Default: 20 seconds
Mode	Enables or disables the server. Default: enabled

In 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 45. Select the **Mode** checkbox to activate the authentication server.
5. Click **Apply** to apply the configuration.



Note: The configuration does not take effect until you perform this step.

In the CLI

```
aaa authentication-server ldap <name>
  host <ipaddr>
  (enter parameters as described in Table 45)
  enable
```

Configuring a TACACS+ Server

Table 46 defines the TACACS+ server parameters.

Table 46 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

In 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 46. Select the **Mode** checkbox to activate the authentication server.
5. Click **Apply** to apply the configuration.



Note: The configuration does not take effect until you perform this step.

In the CLI

```
aaa authentication-server tacacs <name>
  host <ipaddr>
  key <key>
  enable
```

Configuring a Windows Server

Table 47 defines parameters for a Windows server used for stateful NTLM authentication.

Table 47 Windows Server Configuration Parameters

Parameter	Description
Host	IP address of the server. Default: N/A
Mode	Enables or disables the server. Default: enabled

In 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 47](#).
5. Select the **Mode** checkbox to activate the authentication server.
6. Click **Apply** to apply the configuration.



Note: The configuration does not take effect until you perform this step.

In the CLI

```
aaa authentication-server windows <windows-server-name>
  host <ipaddr>
  enable
```

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 in incoming authentication requests is checked against the internal database.

Configuring the Internal Database

By default, the internal database in the master controller is used for authentication. 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 48](#) defines the required and optional parameters used in the internal database.

Table 48 *Internal Database Configuration Parameters*

Parameters	Description
User Name	(Required) Enter a user name or select Generate to automatically generate a user name. An entered username 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. In order 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 244 . (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.

Table 48 Internal Database Configuration Parameters (Continued)

Parameters	Description
Expiration	Select one of the following options: <ul style="list-style-type: none">● Entry does not expire: No expiration on user entry● Set Expiry time (mins): Enter the number of minutes the user will be 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.

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. Enter the information for the client, as described in the table above.
5. Click **Enabled** to activate this entry on creation.
6. Click **Apply** to apply the configuration. The configuration does not take effect until you perform this step
7. At the Servers page, click **Apply**.



Note: 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 518](#).

In the CLI

Enter the following command in enable mode:

```
local-userdb add {generate-username|username <name>} {generate-password|password <password>}L
```

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 *IP address* parameter is added to the existing configuration commands: `local-userdb add`, `local-userdb modify`, `local-userdb-ap add`, and `local-userdb-ap modify`.

In the WebUI

To view *IP address* parameter in the local database, navigate to the **Configuration > Security > Authentication > Servers > Internal DB** page.

Figure 41 IP-Address parameter in the local database

Users									
User Name	Password	Role	E-mail	Enabled	Expiry	IP-Address	Action		
Test	*****	authenticated		Yes		1.2.3.4	Disable	Delete	Modify
rapuser	*****	ap-role		Yes		0.0.0.0	Disable	Delete	Modify
sidd	*****	guest		Yes		0.0.0.0	Disable	Delete	Modify
guest-6381449	*****	guest		Yes		0.0.0.0	Disable	Delete	Modify

To view IP-address parameter in the RAP Whitelist, navigate to the **Wireless > AP Installation > RAP Whitelist** page.

Figure 42 IP-Address parameter in the RAP Whitelist

<input type="checkbox"/>	AP MAC Address	User Name	AP Group	AP Name	Description	Revoked	IP-Address
<input type="checkbox"/>	00:1e:1e:c0:b3:54	q	CP_TEST	AP-125-Port-2	wq		0.0.0.0
<input type="checkbox"/>	00:0b:06:66:02:a9	sid-user	CP_TEST	AP-rap5-port-18	desc		0.0.0.0
<input type="checkbox"/>	11:22:33:44:55:aa	test	grp-4whitelist	testing	test desc		0.0.0.0
<input type="checkbox"/>	ab:cd:ef:11:22:33	aa	New-Group12	12	123		0.0.0.0

New



Note: You cannot configure the IP-Address parameter by using the WebUI.

In 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 add {mac-address <address>} {ap-group|<ap_grup>} {remote-ip<remote-ip>}
local-userdb-ap modify { mac-address <address>} {remote-ip<remote-ip>}
```

The output of **show local-userdb** command:

```
(host) #show local-userdb
```

User Summary

Name	Pwd	Role	E-Mail	Enabled	Expiry	Status	Sponsor-Name	Remote-IP	Grantor-Name
------	-----	------	--------	---------	--------	--------	--------------	-----------	--------------

The output of **show local-userdb-ap** command:

```
(host) #show local-userdb-ap
```

AP-entry Details

Name	AP-Group	AP-Name	Full-Name	Auth-Uname	Rvok-txt	AP_Auth	Descrp	Date-Added	En	Rem-IP
------	----------	---------	-----------	------------	----------	---------	--------	------------	----	--------

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**.

In the CLI

Enter the following command in enable mode:

```
local-userdb export <filename>
local-userdb import <filename>
```

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 User

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 will also remove 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**.

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. The server must be configured before you can include it in a server group.

In 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 menu and click **Add Server**.
 - b. Repeat the above step to add other servers to the group.
6. Click **Apply**.

In the CLI

```
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. By default, 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 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. Dell recommends that you use server selection based on domain matching whenever possible (see [“Configuring Dynamic Server Selection” on page 241](#)).
- 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 fail-through authentication is enabled, users that fail authentication on the first server in the server list should be authenticated with the second server.

In 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](#) 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 menu and click **Add Server**.
12. Repeat [step 11](#) to add ldap-2 to the group.
13. Click **Apply**.

In the CLI

```

aaa authentication-server ldap ldap-1
    host 10.1.1.234
aaa authentication-server ldap ldap-2
    host 10.2.2.234

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 43](#) 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 43 Domain-Based Server Selection Example



You configure the following rules for servers in the corp-serv server group:

- radius-1 will be selected if the client information starts with “host/”.
- radius-2 will be selected if the client information contains “abc.corpnet.com”.

In 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 menu.
 - 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 menu.
 - 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**.



Note: 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**.

In the CLI

```
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.

In 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 menu.
 - 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**.

In the CLI

```
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



Note: This option does not support client information sent in the format host/<pc-name>.<domain>

In 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 menu, then select Trim FQDN, scroll right, and click **Add Server**.
6. Click **Apply**.

In the CLI

```
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.



Note: 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 at <http://technet2.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 49 describes the server rule parameters you can configure.

Table 49 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 <i>Operation</i> and <i>Operand</i> match.
Operation	This is the match method by which the string in <i>Operand</i> is matched with the attribute value returned by the authentication server. <ul style="list-style-type: none"> • contains – The rule is applied if and only if the attribute value contains the string in parameter <i>Operand</i>. • starts-with – The rule is applied if and only if the attribute value returned starts with the string in parameter <i>Operand</i>. • ends-with – The rule is applied if and only if the attribute value returned ends with the string in parameter <i>Operand</i>. • equals – The rule is applied if and only if the attribute value returned equals the string in parameter <i>Operand</i>. • not-equals – The rule is applied if and only if the attribute value returned is not equal to the string in parameter <i>Operand</i>. • 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.

Table 49 Server Rule Configuration Parameters (Continued)

Parameter	Description
Value	The user role or the VLAN applied to the client when the rule is matched.
position	Position of the condition rule. Rules are applied based on the first match principle. 1 is the top. Default: bottom

In 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 menu 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 menu.
 - c. Enter the operand.
 - d. Select **Set VLAN** or **Set Role** from the drop-down menu.
 - e. Enter the value (either user role or VLAN) to be assigned.
 - f. Click **Add**.
 - g. Repeat the above steps to add other rules for the server group.
7. Click **Apply**.

In the CLI

```

aaa server-group <name>
  auth-server <name>
  set {role|vlan} condition <condition> set-value {<role>|<vlan>}
  [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 [“Internal Database” on page 236](#)). In order for 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:

In 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 menu.

- c. Select Set Role from the drop-down menu.
 - d. Click Add.
5. Click Apply.

In the CLI

```
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 50](#)). Accounting is only supported with RADIUS and TACACS+ servers when RADIUS or TACACS+ is used for authentication.

Table 50 *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, see the configuration chapter for the authentication method.

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.



Note: 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.

In the WebUI

1. Navigate to the **Configuration > Management > Administration** page.
2. Under the Management Authentication Servers section, select the Server Group.
3. Click **Apply**.

In the CLI

```
aaa authentication mgmt
    server-group <group>
```

Accounting

You can configure accounting for RADIUS and TACACS+ server groups.



Note: 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. RADIUS accounting works as follows:

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 acknowledgement 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. Currently, possible values include 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

You can use either the WebUI or CLI to assign a server group for RADIUS accounting.

In the WebUI

1. Navigate to the **Configuration > Security > Authentication > AAA Profiles** page.
2. Select AAA Profile, then select the AAA profile instance.
3. Scroll down and select Radius Accounting Server Group. Select the server group from the drop-down menu.
You can add additional servers to the group or configure server rules.
4. Click **Apply**.

In the CLI

```
aaa profile <profile>
    radius-accounting <group>
```

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:

```
aaa tacacs-accounting server-group <group> command {action|all|configuration|show} mode
{enable|disable}
```

Configuring Authentication Timers

[Table 51](#) 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 51 *Authentication Timers*

Timer	Description
User Idle Timeout	<p>Maximum period, in minutes or seconds, after which a client is considered idle if there is no user traffic from the client.</p> <p>The timeout period is reset if there is a user traffic. After this timeout period has elapsed, the controller sends probe packets to the client; if the client responds to the probe, it is considered active and the User Idle Timeout is reset (an active client that is not initiating new sessions is not removed). If the client does not respond to the probe, it is removed from the system. To prevent clients from timing out, set the value in the field to 0. After entering the value, select either min for minutes or sec for seconds from the drop-down list.</p> <p>Range: 0–255 Default: 5 minutes</p>
Authentication Server Dead Time	<p>Maximum period, in minutes, that the controller considers an unresponsive authentication server to be “out of service”.</p> <p>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.</p> <p>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.</p> <p>Range: 0–50 Default: 10 minutes</p>
Logon User Lifetime	<p>Maximum time, in minutes, unauthenticated clients are allowed to remain logged on.</p> <p>Range: 0–255 Default: 5 minutes</p>

Setting an Authentication Timer

To set an authentication timer, complete one of the following procedures:

In 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. Failure to do this results in loss of configuration and you will have to reconfigure the settings.

In the CLI

```
aaa timers {dead-time <minutes>|idle-timeout <number>|logon-lifetime <minutes>}
```

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:

- “Overview of 802.1x Authentication” on page 251
- “Configuring 802.1x Authentication” on page 254
- “Example Configurations” on page 262
- “Advanced Configuration Options for 802.1x” on page 278

Other types of authentication not discussed in this chapter can be found in the following sections of this guide:

- Captive portal authentication: “Captive Portal Authentication” on page 312
- VPN authentication: “VPN Configuration” on page 347
- MAC authentication: “Configuring MAC-Based Authentication” on page 369
- Stateful 802.1x, stateful NTLM, and WISPr authentication: “Stateful and WISPr Authentication” on page 293

Overview of 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 users as well as wireless users.
- The *authenticator* is the gatekeeper to the network and permits or denies access to the supplicants.
- The *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 [http://technet.microsoft.com/en-us/library/cc759077\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc759077(WS.10).aspx)).

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


The 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 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 ensuring the user credentials are kept secure.
- **EAP-GTC**—The EAP-GTC (Generic Token Card) type uses clear text method to exchange authentication controls between client and 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-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 a EAP message. For example, 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 actually 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 client and RADIUS server.
- **ZLXEAP**—This is Zonelabs EAP. For more information, visit <http://tools.ietf.org/html/draft-bersani-eap-synthesis-sharedkeymethods-00#page-30>.

Authentication with a RADIUS Server

See [Table 52](#) 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 44 802.1x Authentication with RADIUS Server



The supplicant and 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.



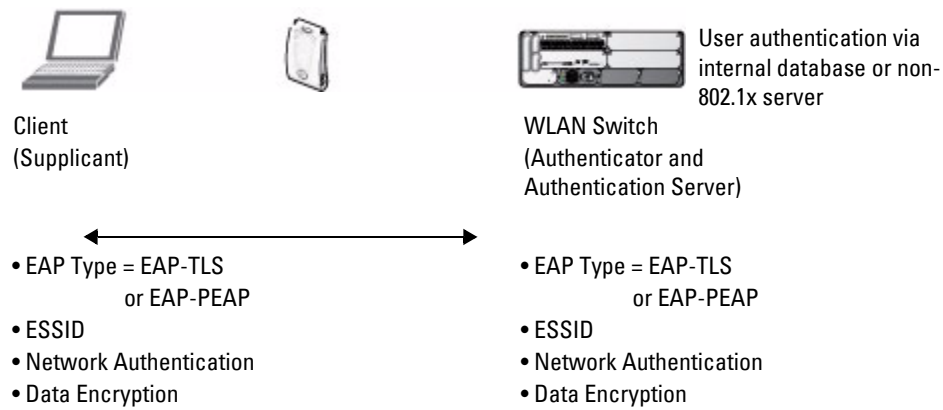
Note: Additional information on EAP types supported in a Windows environment, Microsoft supplicants, and authentication server, is available at [http://technet.microsoft.com/en-us/library/cc782851\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc782851(WS.10).aspx).

The client communicates with the controller through a GRE tunnel in order to form an association with an AP and to authenticate to the network. Therefore, the network authentication and encryption configured for an ESSID must be the same on both the client and the controller.

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 255 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 45 802.1x 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 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 260](#)). The client certificate is verified on the controller (the client certificate must be signed by a known CA) before the user name 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 are using the controller’s internal database for user authentication, you need to add the names and passwords of the users to be authenticated. If you are using an LDAP server for user authentication, you need to configure the LDAP server on the controller, and configure user IDs and passwords. If you are using 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 [Chapter 2, “Network Parameters”](#)
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 [Chapter 10, “Roles and Policies”](#).



Note: 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 and provides differentiated access for different classes of users. For information about obtaining and installing licenses, see [Chapter 27, “Software Licenses”](#).

3. Configure the authentication server(s) and server group. The server can be an 802.1x RADIUS server or, if you are using AAA FastConnect, a non-802.1x server or the controller’s internal database. If you are using EAP-GTC within a PEAP tunnel, you can configure an LDAP or RADIUS server as the authentication server (see [Chapter 8, “Authentication Servers”](#)) If you are using EAP-TLS, you need to import server and CA certificates on the controller (see [“Configuring and Using Certificates with AAA FastConnect” on page 260](#)).
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 [“Using the WebUI” on page 273](#)
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 [“Example Configurations” on page 262](#)

Using 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. The profile details window includes **Basic** and **Advanced** tabs for basic and advanced configuration settings. Click on one or both of these tab to configure the 802.1x Authentication settings. [Table 52](#) describes the parameters you can configure in the high-throughput radio profile.

Table 52 802.1x Authentication Profile Basic WebUI Parameters

Parameter	Description
Basic 802.1x Authentication Profile settings	
Max authentication failures	Number of times a user can try to login with wrong credentials after which the user will be 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. Default: 0
Enforce Machine Authentication	(For Windows environments only) Select this option to enforce machine authentication before user authentication. If selected, either the Machine Authentication Default Role or the User Authentication Default Role is assigned to the user, depending on which authentication is successful. This option is disabled by default. NOTE: This option may require a license (see Chapter 27 on page 527).The Enforce Machine Authentication checkbox is also available on the Advanced settings tab.

Table 52 802.1x Authentication Profile Basic WebUI Parameters (Continued)

Parameter	Description
Machine Authentication: Default Machine Role	Select the default role to be assigned to the user after completing only machine authentication. Default: guest
Machine Authentication: Default User Role	Select the default role to be assigned to the user after completing 802.1x authentication. Default: guest
Reauthentication	Select this option to force the client to do a 802.1x re-authentication after the expiration of the default timer for re-authentication. The default value of the timer (Reauthentication Interval) is 24 hours. If the user fails to re-authenticate with valid credentials, the state of the user is cleared. If derivation rules are used to classify 802.1x-authenticated users, then the Re-authentication timer per role overrides this setting. Default: disabled
Termination	Select this option to terminate 802.1x authentication on the controller. Default: disabled
Termination EAP-Type	The EAP method, either EAP-PEAP or EAP-TLS. Default: eap-peap
Termination Inner EAP-Type	Select one of the following: <ul style="list-style-type: none"> 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 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-Microsoft Challenge Authentication Protocol version 2 (MS-CHAPv2): Described in RFC 2759, this EAP method is widely supported by Microsoft clients. Default: eap-mschapv2
Advanced 802.1x Authentication Profile settings	
Max authentication failures	Number of times a user can try to login 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. The range of allowed values is 0-5 failures, and the default value is 0 failures. NOTE: This option may require a license (see Chapter 27 on page 527).
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 (see Chapter 27 on page 527).
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 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
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.
Interval between Identity Requests	Interval, in seconds, between identity request retries. The allowed range of values is 1-65535 seconds, and the default value is 30 seconds.
Quiet Period after Failed Authentication	The enforced quiet period interval, in seconds, following failed authentication. The allowed range of values is 1-65535 seconds, and the default value is 30 seconds.
Reauthentication Interval	Interval, in seconds, between reauthentication attempts. The allowed range of values for this parameter is 60-864000 seconds, and the default value is 86400 seconds (1day).

Table 52 802.1x Authentication Profile Basic WebUI Parameters (Continued)

Parameter	Description
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. The allowed range of values for this parameter is 60-864000 seconds, and the default value is 1800 seconds.
Unicast Key Rotation Time Interval	Interval, in seconds, between unicast key rotation. The allowed range of values for this parameter is 60-864000 seconds, and the default value is 900 seconds.
Authentication Server Retry Interval	Server group retry interval, in seconds. The allowed range of values for this parameter is 5-65535 seconds, and the default value is 30 seconds.
Authentication Server Retry Count	Maximum number of authentication requests that are sent to server group. The allowed range of values for this parameter is 0-3 requests, and the default value is 2 requests.
Framed MTU	Sets the framed Maximum Transmission Unit (MTU) attribute sent to the authentication server. The allowed range of values for this parameter is 500-1500 bytes, and the default value is 1100 bytes.
Number of times ID-Requests are retried	Maximum number of times ID requests are sent to the client. The allowed range of values for this parameter is 1-10 retries, and the default value is 3 retries.
Maximum Number of Reauthentication Attempts	Number of times a user can try to login 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 may require a license (see Chapter 27 on page 527).
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. (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. The allowed range of values is 1-5 retries, and the default value is 3 retries.
Dynamic WEP Key Size	The default dynamic WEP key size is 128 bits, If desired, you can change this parameter to either 40 bits.
Interval between WPA/WPA2 Key Messages	Interval, in milliseconds, between each WPA key exchange.s The allowed range of values is 1000-5000ms, and the default value is 3000 ms.
Delay between EAP-Success and WPA2 Unicast Key Exchange	Interval, in milliseconds, between unicast and multicast key exchanges. The allowed range of values is 0-2000ms, and the default value is 0 ms (no delay).
Delay between WPA/WPA2 Unicast Key and Group Key Exchange	Interval, in milliseconds, between unicast and multicast key exchanges. The allowed range of values is 0-2000ms, and the default value is 0 ms (no delay).
WPA/WPA2 Key Message Retry Count	Number of times WPA/WPA2 key messages are retried. The allowed range of values for this parameter is 1-5 retries, and the default value is 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.

Table 52 802.1x Authentication Profile Basic WebUI Parameters (Continued)

Parameter	Description
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.
Opportunistic Key Caching	By default, the 802.1x authentication profile enables a cached pairwise master key (PMK) derived via a client and an associated AP and 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 key used by the client.
Validate PMKID	If opp-key-caching is enabled, this option instructs the controller to check the pairwise master key (PMK) ID sent by the client. When this option is enabled, the client must send a PMKID in the associate or reassociate frame to indicate that it supports OKC; otherwise, full 802.1x authentication takes place. (This feature is optional and is disabled by default, since most clients that support OKC do not send the PMKID in their association request.)
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. The range of allowed values is 1024-1500 bytes, and 1300 bytes
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 termination is enabled, click either EAP-PEAP or EAP-TLS to select a Extensible Authentication Protocol (EAP) method.
Termination Inner EAP-Type	If you are using EAP-PEAP as the EAP method, specify one of the following inner EAP types: <ul style="list-style-type: none"> ● 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.
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 not available, 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.

Table 52 802.1x Authentication Profile Basic WebUI Parameters (Continued)

Parameter	Description
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. NOTE: This option may require a license (see Chapter 27 on page 527).
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 (see Chapter 27 on page 527).

7. Click **Apply**.

Using the CLI

The following command configures settings for an 802.1x authentication profiles. Individual parameters are described in [Table 52](#), above.

```

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>}|{wpa-groupkey-
  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 503](#)
- **Client Certificates**—Client certificates are verified on the controller (the client certificate must be signed by a known CA) before the user name 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 505](#)):
 - Controller’s server certificate
 - CA certificate for the CA that signed the client certificates

Using 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 menu 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**.

Using the CLI

```
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 52 on page 255](#)). This tightens the authentication process further since both the device and user need to be authenticated.

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 policies 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 53](#) describes role assignment based on the results of the machine and user authentications.

Table 53 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 fails (for example, the machine information is not present on the server) and user authentication succeeds. Server-derived roles do not apply.	Machine authentication default user role configured in the 802.1x authentication profile.
Passed	Failed	Machine authentication succeeds 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 <i>only</i> 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 assignments would be 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.

VLAN Assignment with Machine Authentication Enabled

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 “About VLAN Assignments” on page 60). 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.



Note: 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 54 describes VLAN assignment based on the results of the machine and user authentications when VLAN derivation is used.

Table 54 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 fails (for example, the machine information is not present on the server) and user authentication succeeds.	VLAN configured in the virtual AP profile
Passed	Failed	Machine authentication succeeds 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.

Example Configurations

The following examples show basic configurations on the controller for:

- “Authentication with an 802.1x RADIUS Server” on page 319
- “Authentication with the Controller’s Internal Database” on page 333

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.

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.



Note: [Appendix D, “802.1x Configuration for IAS and Windows Clients”](#) 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
- Sysadming
- 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.

Using 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**.



Note: 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**. 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**.

Using the CLI

```
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

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**.

Using the CLI

```
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

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.

Using 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 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**.

Using the CLI

```
time-range working-hours periodic
  weekday 07:30 to 17:00

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

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.

Using the WebUI to create the sysadmin role

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**.

Using the CLI to create the sysadmin role

```
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

```
user-role computer
  session-acl allowall
```

Creating an alias for the internal network using CLI

```
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.

Using 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 menu, 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 menu.
 - c. For Operand, select **set role**.
 - d. Click **Add**.
5. Click **Apply**.

Using the CLI

```
aaa authentication-server radius IAS1
    host 10.1.1.21
    key |*a^t%183923!

aaa server-group IAS
    auth-server IAS1
    set role condition Class value-of
```

Configure 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 without machine authentication taking place first, the user is placed in the limited guest role.

Using the WebUI

1. Navigate to the **Configuration > Security > Authentication > L2 Authentication** page.
2. Select 802.1x Authentication Profile.
 - a. In the list of instances, 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 menu, 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 menu, select the IAS server group you created previously.
 - b. Click **Apply**.

Using the CLI

```
aaa authentication dot1x dot1x
    machine-authentication enable
    machine-authentication machine-default-role computer
    machine-authentication user-default-role guest

aaa profile aaa_dot1x
    dot1x-default-role faculty
```



```
mac-default-role computer
authentication-dot1x dot1x
dot1x-server-group IAS
```

Configure 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.

Using 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 > 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**.

Using the CLI

```
vlan 60
interface vlan 60
  ip address 10.1.60.1 255.255.255.0
  ip helper-address 10.1.1.25

vlan 61
interface vlan 61
```

```

ip address 10.1.61.1 255.255.255.0
ip helper-address 10.1.1.25

vlan 63
interface vlan 63
ip address 10.1.63.1 255.255.255.0
ip helper-address 10.1.1.25

ip default-gateway 10.1.1.254

```

Configuring the 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 “[AP Groups](#)” on page 103 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.

Using 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, then select Virtual AP.
4. To create the guest virtual AP:
 - a. Select **NEW** from the Add a profile drop-down menu. 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 menu. 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. Make sure Virtual AP enable is selected.
 - 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, then select Virtual AP.

9. Select **guest** from the Add a profile drop-down menu. Click **Add**.
10. Click **Apply**.

Using the CLI

```
wlan ssid-profile guest
  essid guest
  wepkey1 aaaaaaaaaa
  opmode static-wep

wlan virtual-ap guest
  vlan 63
  ssid-profile guest

ap-group first-floor
  virtual-ap guest
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”.

Using 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**, then select **Virtual AP**.
4. To configure the WLAN-01_first-floor virtual AP:
 - a. Select **NEW** from the Add a profile drop-down menu. 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** in the pop-up window.
 - c. From the SSID profile drop-down menu, 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** in the pop-up window.
 - 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. Make sure **Virtual AP enable** is selected.
 - 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**, then select **Virtual AP**.

9. To configure the WLAN-01_second-floor virtual AP:
 - a. Select **NEW** from the Add a profile drop-down menu. 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 menu. A pop-up window displays the configured AAA profile parameters. Click **Apply** in the pop-up window.
 - c. From the SSID profile drop-down menu, select **WLAN-01**. A pop-up window displays the configured SSID profile parameters. Click **Apply** in the pop-up window.
 - 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. Make sure Virtual AP enable is selected.
 - b. For VLAN, select 61.
 - c. Click **Apply**.

Using the CLI

```
wlan ssid-profile WLAN-01
  essid WLAN-01
  opmode wpa-tkip

wlan virtual-ap WLAN-01_first-floor
  vlan 60
  aaa-profile aaa_dot1x
  ssid-profile WLAN-01

wlan virtual-ap WLAN-01_second-floor
  vlan 61
  aaa-profile aaa_dot1x
  ssid-profile WLAN-01

ap-group first-floor
  virtual-ap WLAN-01_first-floor
ap-group second-floor
  virtual-ap WLAN-01_second-floor
```

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.

Using 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 the 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**.

Using the CLI



Note: Use the privileged mode in the CLI to configure users in the controller's internal database.

```
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 menu.
 - c. Select **Set Role** from the drop-down menu.
 - d. Click **Add**.
5. Click **Apply**.

Configuring a server rule using the CLI

```
aaa server-group internal
  set role condition Role value-of
```

Configure 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.

Using 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 menu.
 - 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**.

Using the CLI

```
aaa authentication dot1x dot1x
    termination enable
```

```
aaa profile aaa_dot1x
    dot1x-default-role student
    authentication-dot1x dot1x
    dot1x-server-group internal
```

Configure 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.

Using 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**.

Using the CLI

```
vlan 60
interface vlan 60
  ip address 10.1.60.1 255.255.255.0
  ip helper-address 10.1.1.25

vlan 61
interface vlan 61
  ip address 10.1.61.1 255.255.255.0
  ip helper-address 10.1.1.25

vlan 63
interface vlan 63
  ip address 10.1.63.1 255.255.255.0
  ip helper-address 10.1.1.25

ip default-gateway 10.1.1.254
```

Configuring the 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 “[AP Groups](#)” on page 103 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.

Using 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 guest virtual AP:
 - a. Select **NEW** from the Add a profile drop-down menu. 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 menu. 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. Make sure Virtual AP enable is selected.
 - 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**, then select Virtual AP.
 9. Select **guest** from the Add a profile drop-down menu. Click **Add**.
 10. Click **Apply**.

Using the CLI

```
wlan ssid-profile guest
  essid guest
  wepkey1 aaaaaaaaaa
  opmode static-wep

wlan virtual-ap guest
  vlan 63
  ssid-profile guest

ap-group first-floor
  virtual-ap guest
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”.

Using 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 menu. 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 menu. A pop-up window displays the configured AAA parameters. Click **Apply** in the pop-up window.
 - c. From the SSID profile drop-down menu, 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** in the pop-up window.
 - 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. Make sure Virtual AP enable is selected.
 - 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** then select **Virtual AP**.
 9. To create the WLAN-01_second-floor virtual AP:
 - a. Select **NEW** from the Add a profile drop-down menu. 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 menu. A pop-up window displays the configured AAA profile parameters. Click **Apply** in the pop-up window.
 - c. From the SSID profile drop-down menu, select **WLAN-01**. a pop-up window displays the configured SSID profile parameters. Click **Apply** in the pop-up window.
 - 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 configuration parameters.
 - a. Make sure Virtual AP enable is selected.
 - b. For VLAN, select **61**.
 - c. Click **Apply**.

Using the CLI

```
wlan ssid-profile WLAN-01
  essid WLAN-01
  opmode wpa-tkip

wlan virtual-ap WLAN-01_first-floor
  vlan 60
  aaa-profile aaa_dot1x
  ssid-profile WLAN-01

wlan virtual-ap WLAN-01_second-floor
  vlan 61
  aaa-profile aaa_dot1x
  ssid-profile WLAN-01

ap-group first-floor
  virtual-ap WLAN-01_first-floor
ap-group second-floor
  virtual-ap WLAN-01_second-floor
```

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.



Note: 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

Using 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**.

Using the CLI

```
aaa authentication dot1x profile
  reauthentication
  timer reauth-period 6011
  unicast-keyrotation
  timer ukey-rotation-period 1021
  multicast-keyrotation
  timer mkey-rotation-period 1867
```

Every 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.

This chapter describes the following topics:

- “Policies” on page 279
- “Creating a Firewall Policy” on page 280
- “Creating an ACL White List” on page 282
- “User Roles” on page 283
- “User Role Assignments” on page 286
- “Global Firewall Parameters” on page 289



Note: This chapter describes configuring firewall policies and parameters that relate to IPv4 traffic. See [Chapter 28, “IPv6 Client Support” on page 535](#) for information about configuring IPv6 firewall policies and parameters.

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 an 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.



Note: You can apply IPv4 and IPv6 firewall policies to the same user role. See [Chapter 28, “IPv6 Client Support” on page 535](#) for information about configuring IPv6 firewall policies.

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.

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.

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 55](#) describes required and optional parameters for a rule.

Table 55 Firewall Policy Rule Parameters

Field	Description
Source (required)	Source of the traffic, which can be one of the following: <ul style="list-style-type: none">• <i>any</i>: Acts as a wildcard and applies to any source address.• <i>user</i>: This refers to traffic from the wireless client.• <i>host</i>: This refers to traffic from a specific host. When this option is chosen, you must configure the IP address of the host.• <i>network</i>: 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.• <i>alias</i>: 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.

Table 55 Firewall Policy Rule Parameters (Continued)

Field	Description
Service (required)	Type of traffic, which can be one of the following: <ul style="list-style-type: none"> • <i>any</i>: This option specifies that this rule applies to any type of traffic. • <i>tcp</i>: Using this option, you configure a range of TCP port(s) to match for the rule to be applied. • <i>udp</i>: Using this option, you configure a range of UDP port(s) to match for the rule to be applied. • <i>service</i>: 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. • <i>protocol</i>: 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: <ul style="list-style-type: none"> • <i>permit</i>: Permits traffic matching this rule. • <i>drop</i>: Drops packets matching this rule without any notification. • <i>reject</i>: Drops the packet and sends an ICMP notification to the traffic source. • <i>src-nat</i>: 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.) • <i>dst-nat</i>: 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 “<i>captiveportal</i>”. • <i>dual-nat</i>: This option performs both source and destination NAT on packets matching the rule. • <i>redirect to tunnel</i>: 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. • <i>redirect to ESI group</i>: 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.
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 “Voice 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. Click **Add** to create a new policy.
3. Enter web-only for the Policy Name.
4. To configure a firewall policy, select IPv4 Session for Policy Type.
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**.



Note: 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

```
ip access-list session web-only
  any any svc-http permit
  any any svc-https permit
```

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.

Configuring a White List Bandwidth Contract 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. In the **IP Protocol Number** field, enter the number for a protocol used by session traffic.

5. In the **Starting Ports** 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.
6. 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.
7. (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 [“Configuring a Bandwidth Contract in the WebUI” on page 285](#)
8. Click **Done**. The ACL displays on the white list section.
9. To delete an entry, click **Delete** next to the entry you want to delete.
10. Click **Apply** to save changes.

Configuring the White List Bandwidth Contract in the CLI

```
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) #firewall cp {deny|permit} proto <IP protocol number> ports <start port number> <last port number> [bandwidth-contract <name>]
```

To create a whitelist ACL entry that permits traffic using protocol 6 on ports 5000 through 6000 to be forwarded to the controller:

```
(host) (config-fw-cp) #firewall cp permit proto 6 ports 5000 6000
```

To create a a whitelist ACL entry that denies traffic using protocol 2 on port 5000 from being forwarded to the controller:

```
(host) (config-fw-cp) #firewall cp deny proto 2 ports 5000 5000
```

User Roles

This section describes how to create a new user role. When you create a user role, you specify one or more policies for the role.

[Table 56](#) describes the different parameters you can configure for the user role.

Table 56 *User Role Parameters*

Field	Description
Firewall Policies (required)	<p>One or more policies that define the privileges of a wireless client in this role. There are three ways to add a firewall policy to a user role:</p> <ul style="list-style-type: none"> • Choose from configured policies (see “Creating a Firewall Policy” on page 280): Select a policy from the list of configured policies and click the “Done” button to add the policy to the list of policies in the user role. If this policy is to be applied to this user role only for specific AP groups, you can specify the applicable AP group. • Create a new policy from a configured policy: This option can be used to create a new policy that is derived from an existing policy. • Create a new policy: The rules for the policy can be added as explained in “Creating a Firewall Policy” on page 280.
Re-authentication Interval (optional)	<p>Time, in minutes, after which the client is required to reauthenticate. Enter a value between 0-4096. 0 disables reauthentication.</p> <p>Default: 0 (disabled)</p>

Table 56 *User Role Parameters (Continued)*

Field	Description
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 “Bandwidth Contracts” on page 285 .
VPN Dialer (optional)	This assigns a VPN dialer to a user role. For details about VPN dialer, see Chapter 15, “Virtual Private Networks” . 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 Chapter 15, “Virtual Private Networks” . 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 Chapter 15, “Virtual Private Networks” . 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 Chapter 12, “Captive Portal” .
Max Sessions	This configures a maximum number of sessions per user in this role. The default is 65535. You can configure any value between 0-65535.

Creating a User Role

The following example creates the user role ‘web-guest’ and assigns the previously-configured ‘web-only’ policy to this user 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 web-guest for Role Name.
4. Under Firewall Policies, click **Add**. From Choose from Configured Policies, select the ‘web-only’ session policy from the list. You can click **Create** to create and configure a new policy.
5. Click **Done** to add the policy to the user role.



Note: If there are multiple policies for this role, policies can be re-ordered by using the up and down buttons provided for each policy.

6. You can optionally enter configuration values as described in [Table 56](#).
7. Click **Apply** to apply this configuration. The role is not created until the configuration is applied.

After assigning the user role (see [“User Role Assignments” on page 286](#)), you can click the Show Reference button to see the profiles that reference this user role.

To 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.



Note: 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

```
user-role web-guest
access-list session web-only position 1
```

After assigning the user role (see “[User Role Assignments](#)” on page 286), you can use the **show reference user-role <role>** command to see the profiles that reference this user role.

Bandwidth Contracts

You can manage bandwidth utilization by assigning maximum bandwidth rates, or *bandwidth contracts*, to user roles. You can configure bandwidth contracts, in kilobits per second (Kbps) or megabits per second (Mbps), for the following types of traffic:

- from the client to the controller (“upstream” traffic)
- from the controller to the client (“downstream” traffic)

You can assign different bandwidth contracts to upstream and downstream traffic for the same user role. You can also assign a bandwidth contract for only upstream or only downstream traffic for a user role; if there is no bandwidth contract specified for a traffic direction, unlimited bandwidth is allowed.

By default, all users that belong to the same role share a configured bandwidth rate for upstream or downstream traffic. You can optionally apply a bandwidth contract on a *per-user* basis; each user who belongs to the role is allowed the configured bandwidth rate.

For example, if clients are connected to the controller through a DSL line, you may want to restrict the upstream bandwidth rate allowed for *each* user to 128 Kbps. Or, you can limit the *total* downstream bandwidth used by all users in the ‘guest’ role to 128 Mbps. The following example configures a bandwidth rate of 128 Kbps and applies it to upstream traffic for the previously-configured ‘web-guest’ user role on a per-user basis.

Configuring a Bandwidth Contract in the WebUI

In the WebUI, you can first configure a bandwidth contract and then assign it to a user role:

1. Navigate to the **Configuration > Advanced Services > Stateful Firewall > BW Contracts** page.
2. Click **Add** to create a new contract.
3. In the **Contract Name** field, enter **BC512_up**.
4. The **Bandwidth** field allows you to define a bandwidth rate in either kbps or Mbps. For this example, enter **512** in the **Bandwidth** field, then click the drop-down list and select **kbps**.
5. Click **Done**.

Assigning a Bandwidth Contract to a User Role in the WebUI

Now that you have a defined bandwidth contract, you can assign that contract to a user role.

1. Navigate to the **Configuration > Security > Access Control > User Roles** page.
2. Select **Edit** for the web-guest user role.
3. Under **Bandwidth Contract**, select **BC512_up** from the drop-down menu for Upstream.
4. Select **Per User**.
5. Scroll to the bottom of the page, and click **Apply**.

You can also can configure the user role and create the bandwidth contract from the **User Roles** page:

1. Navigate to the **Configuration > Security > Access Control > User Roles** page.
2. Select Edit for the web-guest user role.
3. In the Bandwidth Contract section, click the **Upstream** drop-down list and select **Add New**. The **New Bandwidth Contract** fields appear.
 - a. In the **Name** field, enter **BC512_up**.
 - b. In the **Bandwidth** field, enter **512**.
 - c. Click the **Bandwidth** drop-down list and select **kbps**.
 - d. Click **Done** to add the new contract and assign it to the role. The **New Bandwidth Contract** section closes.
4. In the **Bandwidth Contract** section, select the **Per User** checkbox.
5. Scroll to the bottom of the page, and click **Apply**.

Configuring and Assigning Bandwidth Contracts in the CLI

```
aaa bandwidth-contract BC512_up kbps 512
user-role web-guest
  bw-contract BC512_up per-user upstream
```

User Role Assignments

A client is assigned a user role by one of several methods. A user role assigned by one method may take precedence over a user role assigned by a different method. The methods of assigning user roles are, from lowest to highest precedence:

1. The initial user role for unauthenticated clients is configured in the AAA profile for a virtual AP (see [Chapter 4, “Access Points”](#)).
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 an Dell VSA takes precedence over any other user roles.

The following sections describe the methods of assigning user roles.

User Role in AAA Profile

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

```
aaa profile <profile>
  initial-role <role>
  dot1x-default-role <role>
  mac-default-role <role>
```

For additional information on creating AAA profiles, see [“AAA Profile Parameters” on page 118](#).

User-Derived Role

The user role can be derived from attributes from the client’s association with an AP. 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.



Note: User-derivation rules are executed before the client is authenticated.

[Table 57](#) describes the conditions for which you can specify a user role.

Table 57 *Conditions for User-Derived Role*

Rule Type	Condition	Value
BSSID of AP to which client is associated	One of the following: <ul style="list-style-type: none"> • contains • ends with • equals • does not equal • starts with 	MAC address (xx:xx:xx:xx:xx:xx)
User class identifier (option 77) returned by DHCP server	equals	string
Encryption type used by client	One of the following: <ul style="list-style-type: none"> • equals • does not equal 	<ul style="list-style-type: none"> • Open (no encryption) • WPA/WPA2 AES • WPA-TKIP (static or dynamic) • Dynamic WEP • WPA/WPA2 AES PSK • Static WEP • xSec
ESSID to which the client is associated	One of the following: <ul style="list-style-type: none"> • contains • ends with • equals • does not equal • starts with • value of (does not take <i>string</i>; attribute value is used as role) 	string

Table 57 Conditions for User-Derived Role (Continued)

Rule Type	Condition	Value
Location—AP name of the AP to which the client is associated	One of the following: <ul style="list-style-type: none"> • equals • does not equal 	string
MAC address of the client	One of the following: <ul style="list-style-type: none"> • contains • ends with • equals • does not equal • starts with 	MAC address (xx:xx:xx:xx:xx:xx)

Configuring a User-derived Role

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. (You can select VLAN to configure derivation rules for setting the VLAN assigned to a client.)
5. Configure the condition for the rule by setting the Rule Type, Condition, and Value parameters. See [Table 57](#) 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**.

Configure a User-derived Role in the CLI

```
aaa derivation-rules user <name>
    set role condition <condition> set-value <role> position <number>
```

where *condition* consists of *rule_type condition value* parameters. See [Table 57](#) for descriptions of these parameters.

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” on page 299](#).

In the CLI

To configure the default user role for MAC or 802.1x authentication:

```
aaa profile <profile>
  mac-default-role <role>
  dot1x-default-role <role>
```

To configure the default user role for other authentication methods:

```
aaa authentication captive-portal <profile>
  default-role <role>
aaa authentication stateful-dot1x
  default-role <role>
aaa authentication stateful-ntlm
  default-role <role>
aaa authentication vpn
  default-role <role>
```

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. 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 244](#).

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)”.

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.

Global Firewall Parameters

[Table 58](#) 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 [Chapter 28, “IPv6 Client Support”](#) for information about configuring firewall parameters for IPv6 traffic.

Table 58 IPv4 Firewall Parameters

Parameter	Description
Monitor Ping Attack	Number of ICMP pings per second, which if exceeded, can indicate a denial of service attack. Valid range is 1-255 pings per second. Recommended value is 4. Default: No default

Table 58 IPv4 Firewall Parameters (Continued)

Parameter	Description
Monitor TCP SYN Attack rate	Number of TCP SYN messages per second, which if exceeded, can indicate a denial of service attack. Valid range is 1-255 messages per second. Recommended value is 32. Default: No default
Monitor IP Session Attack	Number of TCP or UDP connection requests per second, which if exceeded, can indicate a denial of service attack. Valid range is 1-255 requests per second. Recommended value is 32. Default: No default
Monitor/Police CP Attack rate (per sec)	Rate of misbehaving user's inbound traffic, which if exceeded, can indicate a denial or service attack. Recommended value is 100 frames per second.
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 an Dell representative. Default: Disabled
Prevent L2 Bridging between Wireless users	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
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, 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 an 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 an Dell representative. Default: Disabled
Disable 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 <i>not</i> 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.

Table 58 IPv4 Firewall Parameters (Continued)

Parameter	Description
Session Mirror Destination	<p>Destination (IP address or port) to which mirrored session packets are sent. This option is used only for troubleshooting or debugging.</p> <p>Packets can be mirrored in multiple ACLs, so only a single copy is mirrored if there is a match within more than one ACL.</p> <p>You can configure the following:</p> <ul style="list-style-type: none"> Ethertype to be mirrored with the Ether-type 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. <p>If you configure both an IP address and a port to receive mirrored packets, the IP address takes precedence.</p> <p>Default: N/A</p>
Session Idle Timeout	<p>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 an Dell representative.</p> <p>Default: 15 seconds</p>
Disable FTP Server	<p>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 an Dell representative.</p> <p>Default: Disabled (FTP server is enabled)</p>
GRE Call ID Processing	<p>Creates a unique state for each PPTP tunnel. You should not enable this option unless instructed to do so by an Dell representative.</p> <p>Default: Disabled</p>
Per-packet Logging	<p>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 an Dell representative, as doing so may create unnecessary overhead on the controller.</p> <p>Default: Disabled (per-session logging is performed)</p>
Broadcast-filter ARP	<p>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.</p> <p>Default: Disabled</p>
Session VOIP Timeout (sec)	<p>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.</p> <p>Default: 300 seconds</p>
Disable Stateful H.323 Processing	<p>Disables stateful H.323 processing.</p> <p>Default: Enabled</p>
Disable Stateful SCCP Processing	<p>Disables stateful SCCP processing.</p> <p>Default: Disabled</p>
Only allow local subnets in user table	<p>Adds only IP addresses, which belong to a local subnet, to the user-table.</p> <p>Default: Disabled</p>
Session mirror IPSEC	<p>Configures session mirroring of all frames that are processed by IPsec. Frames are sent to IP address specified by the session-mirror-destination option.</p> <p>NOTE: Use this option for debugging or troubleshooting only.</p> <p>Default: Disabled</p>
Enforce WMM Voice Priority Matches Flow Content	<p>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.</p> <p>Default: Disabled</p>
Rate limit CP untrusted ucast traffic (Mbps)	<p>Specifies the untrusted unicast traffic rate limit. Range is 1-200 Mbps.</p> <p>Default: 10 Mbps</p>

Table 58 IPv4 Firewall Parameters (Continued)

Parameter	Description
Rate limit CP untrusted mcast traffic (Mbps)	Specifies the untrusted multicast traffic rate limit. Range is 1-200 Mbps. Default: 2 Mbps
Rate limit CP trusted ucast traffic (Mbps)	Specifies the trusted unicast traffic rate limit. Range is 1-200 Mbps. Default: 80 Mbps
Rate limit CP trusted mcast traffic (Mbps)	Specifies the trusted multicast traffic rate limit. Range is 1-200 Mbps. Default: 2 Mbps
Rate limit CP route traffic (Mbps)	Specifies the traffic rate limit that needs ARP requests. Range is 1-200 Mbps. Default: 1 Mbps
Rate limit CP session mirror traffic (Mbps)	Specifies the session mirrored traffic forwarded to the controller. Range is 1-200 Mbps. Default: 1 Mbps
Rate limit CP auth process traffic (Mbps)	Specifies the traffic rate limit that is forwarded to the authentication process. Range is 1-200 Mbps. Default: 1 Mbps

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:

- [“Stateful Authentication Overview” on page 293](#)
- [“WISPr Authentication Overview” on page 293](#)
- [“Important Points to Remember” on page 294](#)
- [“Stateful 802.1x Authentication” on page 294](#)
- [“Stateful NTLM Authentication” on page 295](#)
- [“Configuring WISPr Authentication” on page 296](#)

Stateful Authentication Overview

ArubaOS supports two different types of stateful authentication, stateful 802.1x and stateful NTLM.

- **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 a 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 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 can recognize that the client has been authenticated and assign 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.

WISPr Authentication Overview

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 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 will forward 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 will be authenticated on your hotspot's own ISP, as per their service agreements. Once 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 HTML messages to the controller.

- iPass
- Bongo
- Trustive
- weRoam
- AT&T

Important Points to Remember

Before you can configure a stateful authentication feature, you should have defined a user role you want to assign to the authenticated users, and created 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 279](#)
- [“Configuring a RADIUS Server” on page 232](#)
- [“Configuring a Windows Server” on page 235](#)
- [“Server Groups” on page 239](#)

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 802.1x authentication uses only a single Stateful 802.1x profile. This profile can be enabled or disabled, but you can not configure more than one instance of a Stateful 802.1x profile.

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 will perform 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, ArubaOS will note the deauthentication message from the RADIUS server, and will change 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 232](#)

Configure Authentication via 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 that will be assigned to stateful 802.1x authenticated users.
4. Specify the timeout period for authentication requests, from 1-20 seconds. The default value is 10 seconds.
5. Select the **Mode** checkbox to enable stateful 802.1x authentication.

Configure Authentication via the CLI

Use the following commands 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 of commands associates that server group with the stateful 802.1x authentication profile, then sets the authentication role and timeout period.

```
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
  !

aaa server-group group <server-group>
  auth-server <server-name>
  !

aaa authentication stateful-dot1x
  server-group <server-group>
  default-role <role>
  enable
  timeout <seconds>
```

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 will remain 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.

Configure Authentication via 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 *new* Stateful NTLM Authentication profile, select an existing profile, then click the **Save As** button 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 1-20 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**.

Configure Authentication via the CLI

Use the following commands 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, and the third set of commands associates that server group with the stateful NTLM authentication profile then defines the profile settings.

```
aaa authentication-server windows <windows_server_name>
  host <ipaddr>
  enable
  !

aaa server-group group <server-group>
  auth-server <windows_server_name>
  !

aaa authentication stateful-ntlm
  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 will be 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 web sites (www.iso.org and <http://www.itu.int>.)

Configure WISPr Authentication via 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 the **Save As** button in the right window pane. Enter a name for the new profile in the entry field. at the top of the right window pane.

- Define values for the following parameters

Table 59 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 will have 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 CPU utilization threshold value , the Logon wait maximum wait parameter defines the maximum number of seconds a user will have 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 will use the name of the AP to which the user has associated.

- Click **Apply**.
- In the **Profiles** list, select the **Server Group** entry below the WISPr Authentication profile.
- Click the **Server Group** drop-down list and select the group of RADIUS servers you want to use for WISPr authentication.
- Click **Apply**.



Note: 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

Configure WISPr Authentication via the CLI

Use the following CLI commands to configure WISPr authentication. The first set of commands defines the RADIUS server used for WISPr authentication, the second set adds that server to a server group, and the third set of commands associates that server group with the WISPr authentication profile then defines the profile settings.

```

aaa authentication-server radius <rad_server_name>
  host 172.4.77.214
  key qwERtyuIOp
  enable
  nas-identifier corp_venue1
  !

aaa server-group group <server-group>
  auth-server <radius_server_name>
  !

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>
```

Captive portal is one of the methods of authentication supported by ArubaOS. A captive portal presents a web page which requires action on the part of the user 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 Dell controller. For more information about the VPN dialer, see [Chapter 15, “Virtual Private Networks”](#).

This chapter describes the following topics:

- [“Captive Portal Overview” on page 299](#)
- [“Captive Portal in the Base ArubaOS” on page 300](#)
- [“Captive Portal with the PEFNG License” on page 302](#)
- [“Example Authentication with Captive Portal” on page 305](#)
- [“Guest VLANs” on page 311](#)
- [“Captive Portal Authentication” on page 312](#)
- [“Optional Captive Portal Configurations” on page 316](#)
- [“Personalizing the Captive Portal Page” on page 320](#)

Captive Portal Overview

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.



Note: 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 320](#))

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. Later sections in 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 503 in Chapter 26, “Management Access”](#).

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:

```
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:

```
web-server
  captive-portal-cert ServerCert1
  no captive-portal-cert
  captive-portal-cert ServerCert2
```

Captive Portal in the Base ArubaOS

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.

Note: 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 [Chapter 8, “Authentication Servers”](#).
- 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 1](#). 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 1](#).
- 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.

Note: In ArubaOS 2.5.2 and later 2.5.x releases, captive portal users in the base operating system are placed into the predefined *cpbase* initial user role before authentication. The *cpbase* role is not supported in ArubaOS 3.x. You need to create new captive portal profiles in the base operating system, as described in this section, which automatically generates the required policies and roles.

Configuring Captive Portal via the WebUI

1. Navigate to the **Configuration > Security > Authentication > L3 Authentication** page. 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. You can enable user login and/or guest login, and configure other captive portal profile parameters as described in [Table 60](#).
 - 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.



Note: 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**.

Configuring Captive Portal via 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:

```
aaa authentication captive-portal c-portal
  server-group cp-srv
aaa profile aaa_c-portal
  initial-role c-portal
wlan ssid-profile ssid_c-portal
  essid c-portal-ap
wlan virtual-ap vp_c-portal
  aaa-profile aaa_c-portal
  ssid-profile ssid_c-portal
  vlan 20
```

Captive Portal with the 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.



Note: MAC-based authentication, if enabled on the controller, takes precedence over captive portal authentication. If you use captive portal, do not enable MAC-based 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 [Chapter 27, “Software Licenses”](#)).

- Configure the user role for a default user.
Create and configure user roles and policies for guest or registered captive portal users. (See [Chapter 10, “Roles and Policies”](#) 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 [Chapter 8, “Authentication Servers”](#) for more information about configuring authentication servers and server groups.)



Note: 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 [Chapter 8, “Authentication Servers”](#).

- 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 via 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 60](#)).
 - 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 via 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:

```
aaa authentication captive-portal c-portal
  default-role employee
  server-group cp-srv
user-role logon
  captive-portal c-portal
aaa profile aaa_c-portal
  initial-role logon
wlan ssid-profile ssid_c-portal
  essid c-portal-ap
  vlan 20
wlan virtual-ap vp_c-portal
  aaa-profile aaa_c-portal
  ssid-profile ssid_c-portal
```

Example 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.



Note: 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-logon 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.



Note: 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.

Configure Policies and Roles via the WebUI

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 **user**.
 - 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**.



Note: The following step defines an alias representing the public DNS server addresses. Once defined, you can use the alias for other rules and policies.

- 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**, 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**.
8. Click **Apply**.

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 **user**.
 - 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**.

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**.



Note: 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**.

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**.

Guest-logon 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 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**.

Guest-Logon 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 clogout 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**.

Configure Policies and Roles via the CLI

Time Range

To create a time range via the command-line interface, access the CLI in config mode and issue the following commands:

```
time-range working-hours periodic
  weekday 07:30 to 17:00
```

Create Aliases

To create a aliases via the command-line interface, access the CLI in config mode and issue the following commands:

```
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
netdestination "Public DNS"
  host 64.151.103.120
  host 216.87.84.209
```

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:

```
ip access-list session guest-logon-access
  user any udp 68 deny
  user any svc-dhcp permit time-range working-hours
  user alias "Public DNS" svc-dns src-nat time-range working-hours
```

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:

```
ip access-list session auth-guest-access
  user any udp 68 deny
  user 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
```

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:

```
ip access-list session block-internal-access
  user alias "Internal Network" any deny
```

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:

```
ip access-list session drop-and-log
  user any any deny log
```

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:

```
user-role guest-logon
  session-acl captiveportal position 1
  session-acl guest-logon-access position 2
  session-acl block-internal-access position 3
```

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:

```
user-role auth-guest
  session-acl clogout 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
```

Guest VLANs

Guests using the WLAN are assigned to VLAN 900 and are given IP addresses via DHCP from the controller.

Configuring the guest VLAN via the WebUI

1. Navigate to the **Configuration > Network > VLANs** page.
 - a. Click **Add**.
 - b. For **VLAN ID**, enter 900.
 - c. Click **Apply**.
2. Navigate to the **Configuration > Network > IP > IP Interfaces** page.
 - 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. For **Pool Name**, enter **guestpool**.
 - d. For **Default Router**, enter **192.168.200.20**.
 - e. For **DNS Server**, enter **64.151.103.120**.
 - f. For **Lease**, enter **4 hours**.
 - g. For **Network**, enter **192.168.200.0**. For **Netmask**, enter **255.255.255.0**.
 - h. Click **Done**.
4. Click **Apply**.

Configuring the guest VLAN via the CLI

```

vlan 900
interface vlan 900
ip address 192.168.200.20 255.255.255.0
ip dhcp pool "guestpool"
default-router 192.168.200.20
dns-server 64.151.103.120
lease 0 4 0
network 192.168.200.0 255.255.255.0

```

Captive Portal Authentication

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:

```

aaa authentication captive-portal guestnet
default-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:

```
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:

```
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:

```
wlan ssid-profile guestnet
  essid guestnet
  opmode opensystem

aaa profile guestnet
  initial-role guest-logon

wlan virtual-ap guestnet
  vlan 900
  aaa-profile guestnet
  ssid-profile guestnet
```

User Account Administration

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 518](#) for more information about configuring guest provisioning users and administering guest accounts.

Captive Portal Configuration Parameters

[Table 60](#) describes configuration parameters on the WebUI Captive Portal Authentication profile page.



Note: In the CLI, you configure these options with the **aaa authentication captive-portal** commands.

Table 60 *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. The Policy Enforcement Firewall Next Generation (PEFNG) license must be installed. 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.
User Login	Enables Captive Portal with authentication of user credentials. Default: enabled
Guest Login	Enables Captive Portal logon without authentication. Default: disabled
Logout popup 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 captiveportal 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	Maximum 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: 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 FQDN	Allows the user to see and select the fully-qualified domain name (FQDN) on the login page. Default: disabled
Use CHAP	Use CHAP protocol. You should not use this option unless instructed to do so by an Dell representative. Default: PAP
Sygate-on-demand-agent	Enables client remediation with Sygate-on-demand-agent (SODA). Default: disabled
Login page	URL of the page that appears for the user 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	Enables the display of the welcome page. If this option is disabled, redirection to the web URL happens immediately after logon. Default: Enabled
Proxy Server Configuration	Configures IP address and port number for proxy server. NOTE: This option is only available in the base operating system. Default: N/A

Table 60 *Captive Portal Authentication Profile Parameters (Continued)*

Parameter	Description
Adding switch ip address in redirection URL	Select this checkbox to add the switch ip address in the redirection URL. This parameter requires the Public Access license.
Adding user vlan in redirection URL	Select this checkbox to add the user VLAN in the redirection URL Default: disabled This parameter requires the Public Access license.
Add a controller interface in redirection URL	Enter the IP address of a controller interface to add that 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.
Allow only one active user session	Select this checkbox to allow only one active user session at a time. This feature is disabled by default.
Show the acceptable use policy page	Select this checkbox to display the acceptable user policy before the login page.
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. This parameter requires the Public Access license.

Optional Captive Portal Configurations

The following are optional captive portal configurations:

- “Per-SSID Captive Portal Page” on page 316
- “Changing the Protocol to HTTP” on page 317
- “Proxy Server Redirect” on page 318
- “Redirecting Clients on Different VLANs” on page 319
- “Web Client Configuration with Proxy Script” on page 319

Per-SSID Captive Portal Page

You can upload custom login pages for captive portal into the controller through the WebUI (refer to [Appendix E, “Internal Captive Portal”](#)). 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 61](#).

Table 61 *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:

```
aaa authentication captive-portal profile
    protocol-http
```

```
(For captive portal with role-based access only)
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
```

Proxy Server Redirect

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.



Note: 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:

```
aaa authentication captive-portal profile
  proxy host ipaddr port port
```

For captive portal with role-based access:

```
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.



Note: 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:

```
ip cp-redirect-address ipaddr
```

For captive portal with PEFNG license:

```
netdestination cp-redirect ipaddr
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 to 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:

```
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

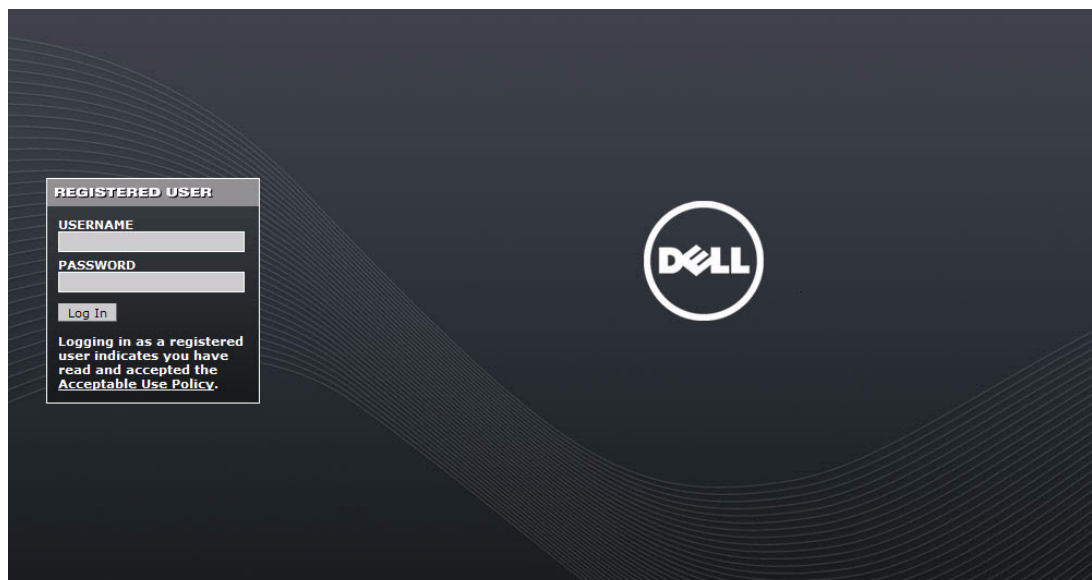
The background image and text should be visible to users with a browser window on a 1024 by 768 pixel screen. The background should not clash if viewed on a much larger monitor. A good option is to have the background image at 800 by 600 pixels, and set the background color to be compatible. The maximum 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 [Appendix E, “Internal Captive Portal”](#)


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.



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.



Note: xSec is an optional ArubaOS software module. You must purchase and install the license for the xSec software module on the controller.

This chapter describes the following topics:

- [“Securing Client Traffic” on page 324](#)
- [“Securing Controller-to-Controller Communication” on page 328](#)
- [“Configuring the Odyssey Client on Client Machines” on page 329](#)

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...



Note: For information about the currently supported release for Funk Odyssey, please contact Juniper Networks.



Note: xSec is an optional licensed feature for Dell controllers. xSec is automatically enabled on the controller when you install the license.

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 [Chapter 2, “Network Parameters”](#) for more information.
2. Configure the user role for the authenticated xSec clients. See [Chapter 10, “Roles and Policies”](#) for information.
3. Configure the server group that will be used to authenticate clients using 802.1x. See [Chapter 8, “Authentication Servers”](#) for more information.
4. Configure the AAA profile to specify the 802.1x default user role. Specify the 802.1x authentication server group.




Note: You can configure the 802.1x authentication profile if necessary. See [Chapter 9, “802.1x Authentication”](#) 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 46](#) 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 46 *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
  dot1x-default-role employee
  dot1x-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 [Chapter 2, “Network Parameters”](#) 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 [Chapter 10, "Roles and Policies"](#) for information.
3. Configure the server group that will be used to authenticate clients using 802.1x. See [Chapter 8, "Authentication Servers"](#) for more information.
4. Configure the controller port to which the wired client(s) 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 47](#) 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 47 *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 47](#), this is VLAN 1.
 - d. For **xSec VLAN**, select the VLAN to which authenticated users are assigned from the drop-down menu. In [Figure 47](#), 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
  dot1x-default-role employee
  dot1x-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 [Chapter 2, “Network Parameters”](#) 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 [Chapter 10, “Roles and Policies”](#) for information.
3. Configure the server group that will be used to authenticate clients using 802.1x. See [Chapter 8, “Authentication Servers”](#) 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
  dot1x-default-role employee
  dot1x-server-group xsec-svrs
aaa authentication wired
  profile xsec-wired
```

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. is 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 48 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 Dell Networks or 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**.

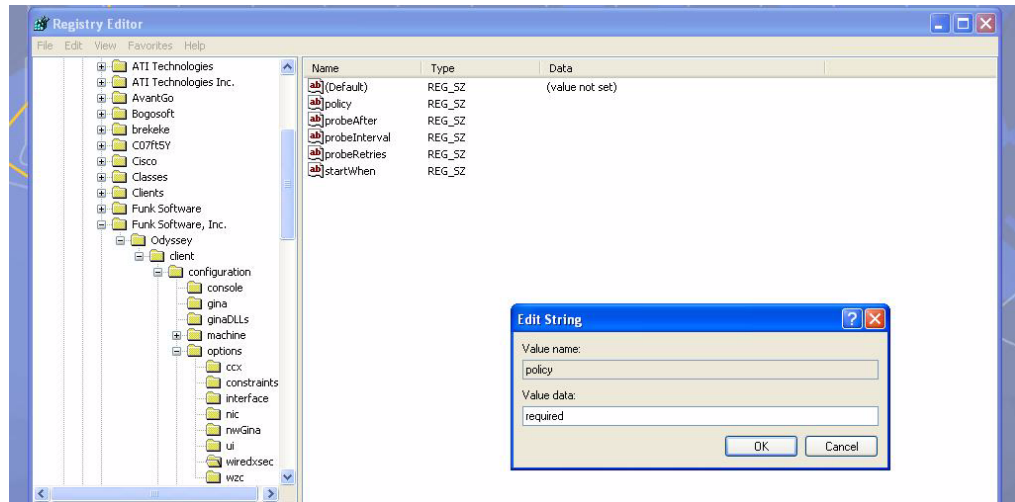
- c. Navigate down the tree to
HKEY_LOCAL_MACHINE\SOFTWARE\Funk Software,
Inc.\odyssey\client\configuration\options\wiredxsec.

Figure 49 *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 50 *Modifying a regedit Policy*



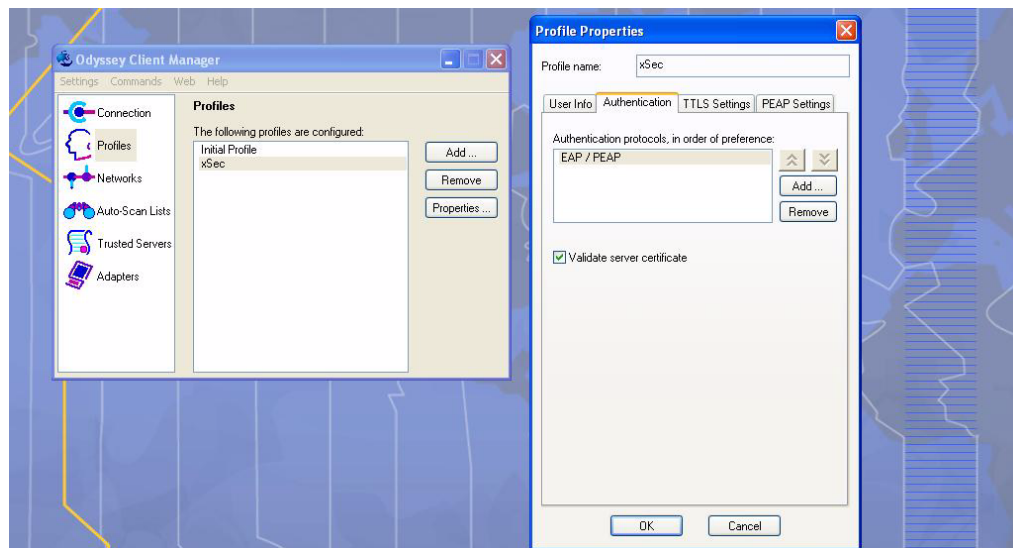
3. Open the Funk Odyssey Client. Click the **Profile** tab in the client window. This allows the user to create the user profile for 802.1x authentication.

Figure 51 *The Funk Odyssey Client Profile*



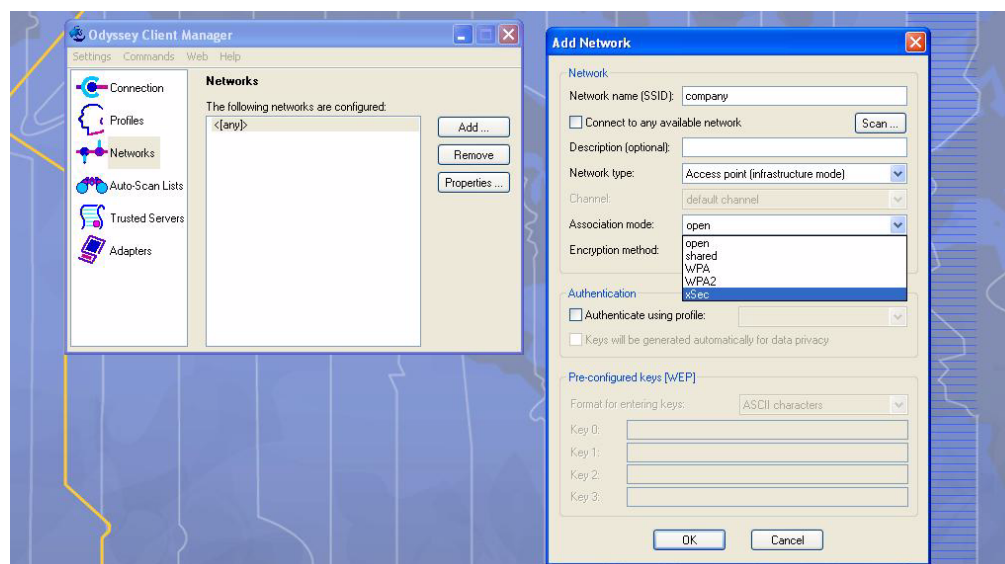
- a. In the login name dialog box, enter the login name used for 802.1x authentication. For the password, the client could use the WINDOWS password or use the configured password based on the selection made.
- b. Click the certificate tab and enter the certificate information required. This example shows the PEAP settings.

Figure 52 *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 53 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.

Virtual Intranet Access (VIA) is part of the Dell remote networks solution targeted for teleworkers and mobile users. It 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 virtual intranet access solution comes in two parts—the VIA (Windows desktop application) and the controller configuration.

- VIA —Teleworkers and mobile users can install a light weight application on their Microsoft Windows computers to connect to their enterprise network from remote locations (see “VIA” on page 333).
- 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 Command Line to configure your controller (see “Configuring the VIA Controller” on page 334).

Topics in this Document

- “VIA” on page 333
- “Configuring the VIA Controller” on page 334

VIA

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.



Note: VIA will be referred to as VIA from this point on in this chapter.

How it Works

VIA provides seamless connectivity experience to users when accessing 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 on system start and establish a remote connection. The following table explains the typical behavior:



Note: The sequence of events described in Table 62 does not necessarily mean that the events always happen in the order shown in the table.

Table 62 VIA Connectivity Behavior

User action / environment	VIA's behavior
The client moves from a trusted to un-trusted environment. <i>Example: From office to a public hot-spot.</i>	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 <i>Settings</i> tab.

Table 62 *VIA Connectivity Behavior*

User action / environment	VIA's behavior
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: <ol style="list-style-type: none"> 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 VIA

Users can download VIA from a URL provided to them by their IT department and install it on their computers. Alternatively, you can install VIA on users computer using SMS (Systems Management Server) or other similar mechanisms.

Upgrade Workflow

VIA checks for upgrade requirements during the login phase. There are two types of upgrade process:

- Minimal Upgrade—This type of upgrade is initiated for bug fixes and some minor enhancements which requires only 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. The user can choose to proceed or cancel the upgrade process. If the user chooses to upgrade, a system reboot is required. If the user cancels the upgrade, VIA will prompt the user for an upgrade every time it attempts to connect.



Note: See [Appendix F, “VIA End User Instructions”](#) for information about using the desktop application.

Configuring the VIA Controller

VIA configuration requires that you first configure VPN settings and then configure VIA settings. See [Chapter 15, “Virtual Private Networks”](#) on page 347 for information on configuring VPN settings on your controller.

Before you Begin

Enable port 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.

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**—The BaseOS 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.
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`
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.
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.
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 ([step 3 on page 335](#)) is configured, users can view this list and select one during the client login.
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.
7. **Configure VIA Client WLAN Profiles**—You can push WLAN profiles to users 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 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.
8. **Re-branding VIA and Downloading the Installer**—You can use a custom logo on the VIA client and on the VIA download web page.
9. **Download VIA Installer and Version File**

Using 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 [Chapter 27, “Software Licenses” on page 527](#) 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.

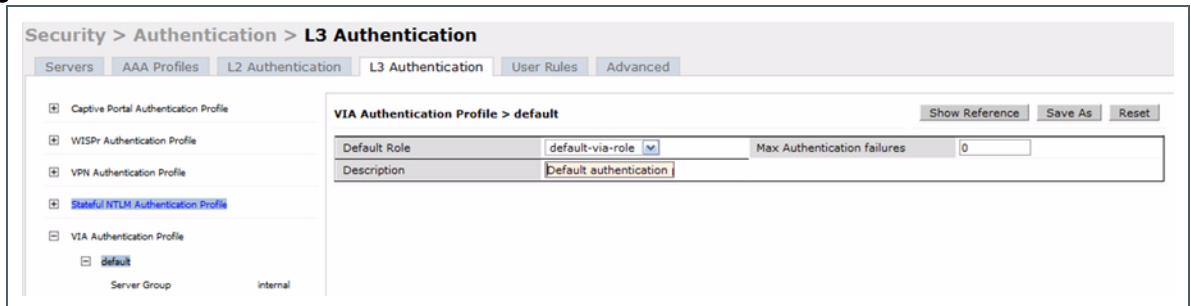
1. Navigate to **Configuration > Security > Authentication > L3 Authentication**.
2. Under the *Profiles* section, expand the **VIA Authentication Profile** option. You can configure the following parameters for the authentication profile:

Table 63 VIA - Authentication Profile Parameters

Parameter	Description
Default Role	This 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.

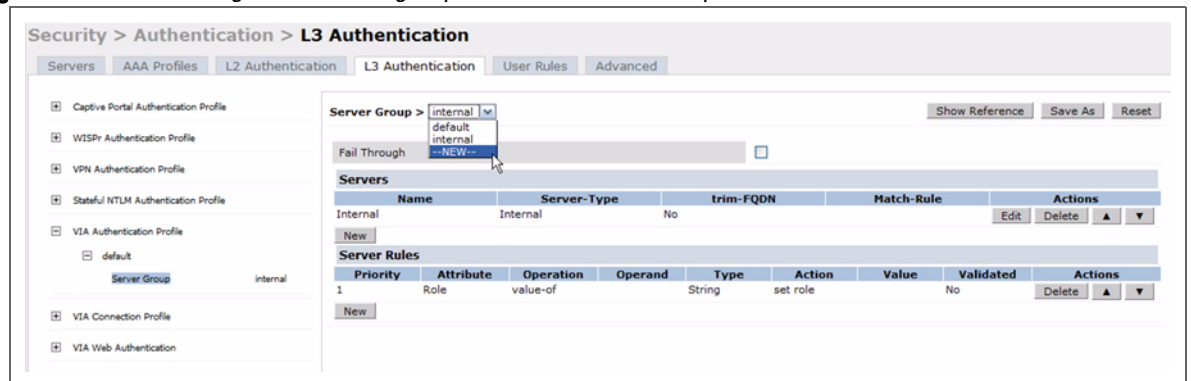
3. To create a new authentication profile:
 1. Enter a name for the new authentication profile under the *VIA Authentication Profiles* section and click the **Add** button.
 2. Expand the **VIA Authentication Profiles** option and select the new profile name.
4. To modify an authentication profile, select the profile name to configure the default role
The following screenshot uses the default authentication profile.

Figure 54 VIA - Associate User Role to VIA Authentication Profile



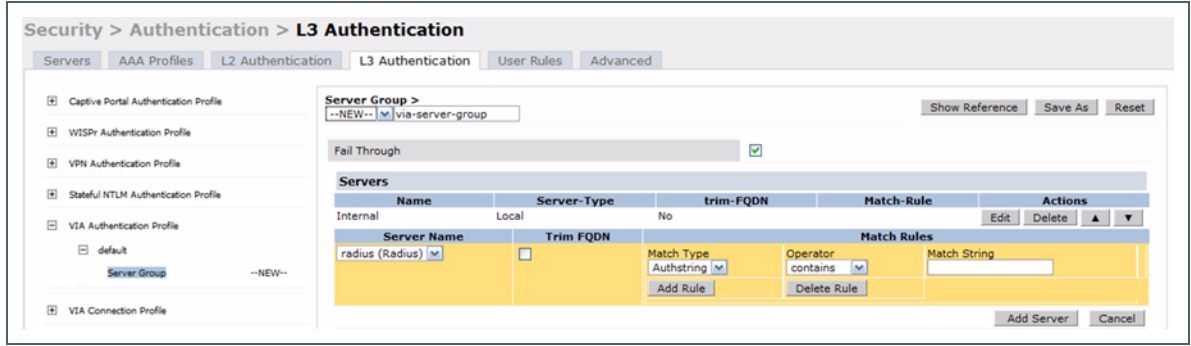
5. To use a different server group, Click **Server Group** under VIA Authentication Profile and select **New** to create a new server group.

Figure 55 VIA - Creating a new server group for VIA authentication profile



6. Enter a name for the server group.

Figure 56 VIA - Enter a name for the server group

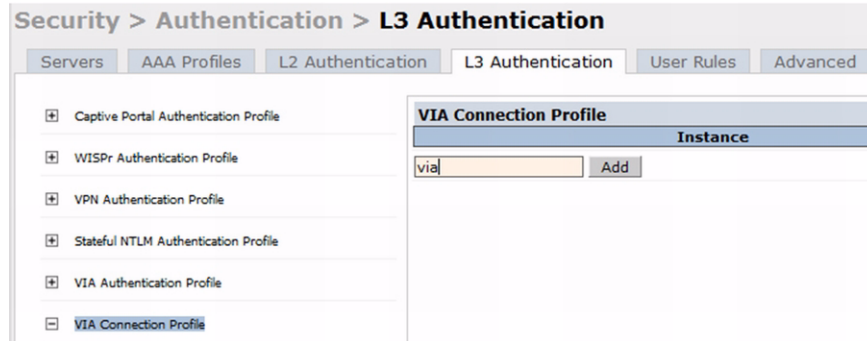


Create VIA Connection Profile

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 57 VIA - Create VIA Connection Profile---redo image



2. Now click on the new VIA connection profile to configure the connection settings:

Figure 58 VIA - Configure VIA Connection Profile

You can configure the following options for a VIA connection profile.

Table 64 VIA - Connection Profile Options

Configuration Option	Description
1	<p>VIA Controller</p> <p>Enter the following information about the VIA controller.</p> <ul style="list-style-type: none"> • 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. <p>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 <i>Up</i> and <i>Down</i> arrows.</p> <p>To delete a controller from your list, select a controller and click the Delete button.</p>
2	<p>VIA Authentication Profiles to provision</p> <p>This is the list of VIA authentication profiles that will be displayed to users in the VIA client. See “Create VIA Authentication Profile” on page 336.</p> <ul style="list-style-type: none"> • 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 <i>Up</i> and <i>Down</i> arrows. • To delete an authentication profile, select a profile name and click the Delete button.
3	<p>VIA tunneled networks</p> <p>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.</p> <ul style="list-style-type: none"> • 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.

Table 64 VIA - Connection Profile Options

	Configuration Option	Description
4	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. <ul style="list-style-type: none"> 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 341 for more information.
5	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 .
6	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 <code>crypto-local ipsec-map <ipsec-map-name></code> command.
7	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
8	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.
9	VIA Support E-mail Address	The support e-mail address to which VIA users will send client logs. Default: None.
10	VIA external download URL	End users will use this URL to download VIA on their computers.
11	Client Auto-Login	Enable or disable VIA client to auto login and establish a secure connection to the controller. Default: Enabled
12	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
13	Enable split-tunneling	Enable or disable split tunneling. <ul style="list-style-type: none"> 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
14	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
15	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
16	Allow user to save passwords	Enable or disable users to save passwords entered in VIA. Default: Enabled
17	Validate Server Certificate	Enable or disable VIA from validating the server certificate presented by the controller. Default: Enabled
18	VIA max session timeout	The maximum time (minutes) allowed before the VIA session is disconnected. Default: 1440 min
19	Maximum reconnection attempts	The maximum number of re-connection attempts by the VIA client due to authentication failures. Default: 3

Table 64 VIA - Connection Profile Options

Configuration Option	Description
20 Allow user to disconnect VIA	Enable or disable users to disconnect their VIA sessions. Default: on

Configure VIA Web Authentication

To configure VIA web authentication profile:

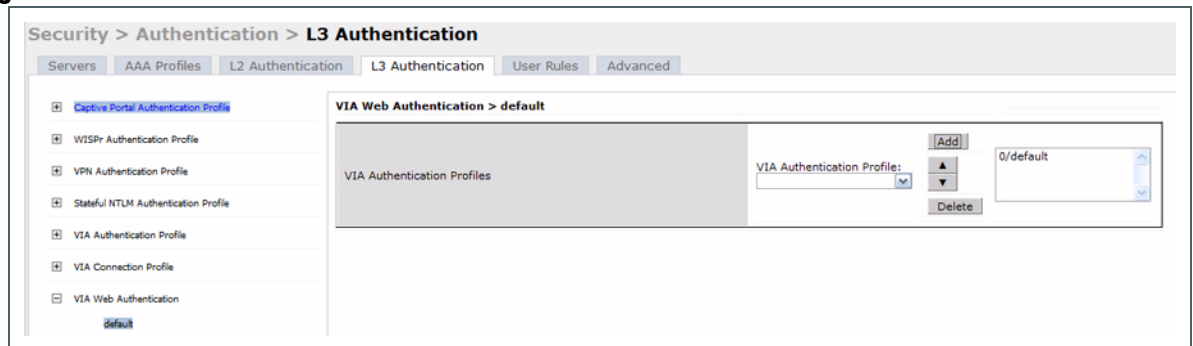
1. Navigate to **Configuration > Security > Authentication > L3 Authentication** tab.
2. Expand VIA Web Authentication and click on *default* profile.



Note: 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 59 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 335](#)) 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 and click the **Change** button.
4. Click the **Apply** button to save the changes to the configuration.

Figure 60 VIA - Associate VIA Connection Profile to User Role

The screenshot shows a configuration page with several sections, each with a dropdown menu and a 'Change' button:

- Re-authentication Interval:** Disabled, Change (0 disables re-authentication. A positive value enables authentication 0 - 4096)
- Role VLAN ID:** Not Assigned, Change
- Bandwidth Contract:** Upstream: Not Enforced, Change, Per User; Downstream: Not Enforced, Change, Per User
- VPN Dialer:** Not Assigned, Change
- L2TP Pool:** default-l2tp-pool, Not Assigned, Change
- PPTP Pool:** default-pptp-pool, Not Assigned, Change
- Captive Portal Profile:** Not Assigned, Change
- VIA Connection Profile:** Not Assigned, Change. The dropdown menu is open, showing 'Not Assigned' and 'Via'.
- Max Sessions:** 65535, Change (0 - 65535)
- Stateful NTLM Profile:** Not Assigned, Change
- WISPr Profile:** Not Assigned, Change

Configure VIA Client WLAN Profiles

To configure a VIA client WLAN profile:

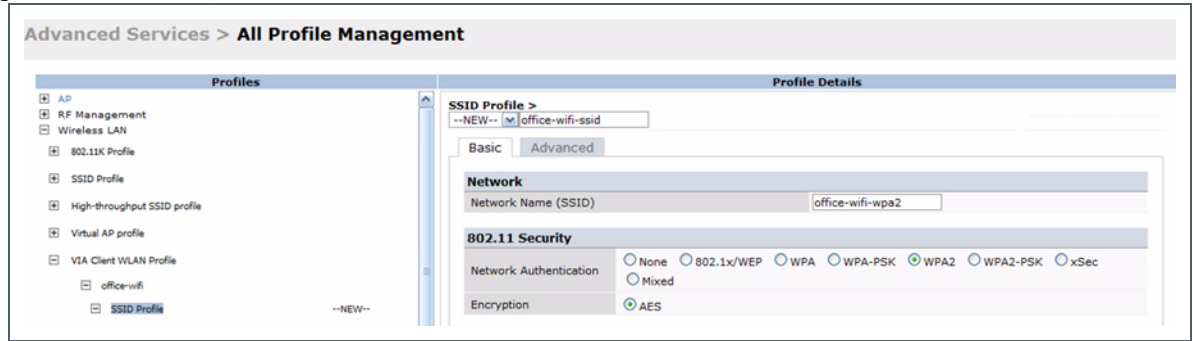
1. Navigate to Configuration > Advanced Services > All Profiles.
2. Expand Controller Profiles and select VIA Client WLAN Profile.
3. In the Profile Details, enter a name for the WLAN profile and click the Add button.

Figure 61 VIA - Create VIA Client WLAN Profile

The screenshot shows the 'Advanced Services > All Profile Management' page. The left sidebar shows a tree view of profiles, with 'VIA Client WLAN Profile' selected. The right pane shows the 'VIA Client WLAN Profile' details, with the 'Instance' field containing the text 'office-wifi' and an 'Add' button.

4. Expand the new WLAN profile and click SSID Profile. In the profile details page, select New from the SSID 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 62 VIA - Configure the SSID Profile



6. You can now configure the SSID profile by select the SSID profile under VIA Client WLAN Profile option.

Figure 63 VIA - Configure VIA Client WLAN Profile

EAP Type	eap-peap		Inner EAP Type	eap-mschapv2	
EAP-PEAP options	<input type="checkbox"/> validate-server-certificate	<input checked="" type="checkbox"/> enable-fast-reconnect	<input type="checkbox"/> enable-quarantine-checks	<input type="checkbox"/> disconnect-if-no-cryptobinding-tlv	
	<input type="checkbox"/> dont-allow-user-authorization				
EAP-Certificate options	<input type="checkbox"/> use-smartcard				
	<input type="checkbox"/> simple-certificate-selection	<input checked="" type="checkbox"/> validate-server-certificate			
	<input type="checkbox"/> use-different-name				
Inner EAP Authentication options	<input type="checkbox"/> mschapv2-use-windows-credentials				
	<input type="checkbox"/> use-smartcard				
	<input type="checkbox"/> simple-certificate-selection	<input checked="" type="checkbox"/> validate-server-certificate			
	<input type="checkbox"/> use-different-name				
Automatically connect when this WLAN is in range	<input checked="" type="checkbox"/>	EAP-PEAP: Connect only to these servers	<input type="text"/>		
Enable IEEE 802.1x authentication for this network	<input checked="" type="checkbox"/>	EAP-Certificate: Connect only to these servers	<input type="text"/>		
Authenticate as computer when computer info is available	<input checked="" type="checkbox"/>	Inner EAP-Certificate: Connect only to these servers	<input type="text"/>		
Authenticate as guest when computer or user info is unavailable	<input type="checkbox"/>	Connect even if this WLAN is not broadcasting	<input type="checkbox"/>		

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 65 Configure VIA client WLAN profile

Option	Description
EAP-PEAP options	<p>Select the following options, if the EAP type is PEAP (Protected EAP):</p> <ul style="list-style-type: none"> 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	<p>Select an EAP type used by client to connect to wireless network. Default: EAP-PEAP</p>

Table 65 Configure VIA client WLAN profile

Option	Description
EAP-Certificate Options	If you select EAP type as certificate, you can select one of the following options: <ul style="list-style-type: none"> ● 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
Inner EAP Authentication options:	<ul style="list-style-type: none"> ● 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

Re-branding VIA and Downloading the Installer

You can re-brand the VIA client, the VIA download page with your custom logo and HTML page.

Figure 64 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 **Configuration > Advanced Services > VPN Services > VIA** tab.
2. Under VIA installers for various platforms section, click **ansetup.msi** to download the installation file. Using CLI to Configure VIA

Customize VIA Logo

To use a custom logo on the VIA download page and the VIA client:

1. Navigate to **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 **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: ``)
- `<% logout %>`: the logout link (Example: `<a href="<% logout %>">VIA Web Logout`)
- `<% download %>`: the installer download link (Example: `<a href="<% download %>">Click here to download VIA`)

To use the default welcome page, click the **Reset** button.

4. Click the **Apply** button to continue.

Using CLI to Configure VIA

The following steps illustrate configuring VIA using CLI.

Enable VPN module

Insert your Policy Enforcement Firewall Next Generation (PEFNG) license key to enable the VPN module on your controller.

```
(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") #controller addr 202.100.10.100 internal-ip 10.11.12.13 desc "VIA Primary Controller" 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 341](#)

```
(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
```



Note: 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 *ArubaOS 5.0 Command Reference Guide*.

Customize VIA logo, landing page and downloading installer

This step can only be performed using the WebUI. See [“Re-branding VIA and Downloading the Installer” on page 343](#).

For wireless networks, virtual private network (VPN) connections can be used to further secure the 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:

- “VPN Configuration” on page 347
- “Remote Access VPN for L2TP IPsec” on page 349
- “Remote Access VPNs for XAuth” on page 355
- “Remote Access VPN for PPTP” on page 362
- “Site-to-Site VPNs” on page 362
- “Dell Dialer” on page 366

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 and 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)
- 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 [Chapter 10, “Roles and Policies”](#) for information about configuring user roles.
- The authentication server group the controller will use to validate the clients. See [Chapter 8, “Authentication Servers”](#) for configuration details.



Note: 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 following sections.

VPN authentication

To configure VPN authentication via the WebUI:

1. Navigate to the **Configuration > Security > Authentication > L3 Authentication** page.
2. In the Profiles list, select the **default** VPN Authentication Profile.
3. Select the Default Role from the drop-down menu.
4. (Optional) Set Max Authentication failures to an integer value (the default value is 0, which disables this feature). This number indicates the number of contiguous authentication failures before the station is blacklisted.
5. Click **Apply**.
6. In the **default** profile list, select Server Group.
7. From the drop-down menu, select the server group to be used for VPN authentication.
8. Click **Apply**.

To configure VPN authentication via the command-line interface, access the CLI in config mode and issue the following commands:

```
aaa authentication vpn default
  default-role <role>
  max-authentication-failure <number>
  server-group <name>
```

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 66](#).

Each row in this table specifies the allowed combinations of AAA servers for simultaneous deployment. Configuration rules include:

- RAP-certs can only use Local-DP-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 irrespective of RAP/CAP authentication configuration server.

Table 66 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

Remote Access VPN for L2TP IPsec

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 as well as 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 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.

Configure the VPN via the WebUI

Authentication Method and Server Addresses

1. Navigate to **Configuration > Advanced Services > VPN Services** and click the **IPsec** tab.
2. To enable L2TP, select **Enable L2TP** (this is enabled by default).
3. Select the authentication method. Currently supported methods are:
 - Password Authentication Protocol (PAP)
 - Extensible Authentication Protocol (EAP)
 - Challenge Handshake Authentication Protocol (CHAP)
 - Microsoft Challenge Handshake Authentication Protocol (MSCHAP)
 - MSCHAP version 2 (MSCHAPv2)
4. 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 will be pushed to the VPN client.

Define 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 start address, the end address and the pool name.
3. Click **Done** to apply the configuration.

Source NAT

1. 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
2. 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:
 - a. Navigate to **Configuration > IP > NAT Pools**.
 - b. Click **Add**.
 - c. In the **Pool Name** field, enter a name for the new NAT pool, up to 63 alphanumeric characters.
 - d. 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.
 - e. 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.

- f. 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 will use the destination NAT IP 0.0.0.0.
- g. Click **Done** to close the NAT pools tab
- h. Navigate to **Configuration > Advanced Services > VPN Services** and click the **IPsec** tab to return to the IPsec window.
- i. Click the **NAT Pool** drop-down list and select the NAT pool you just created.

IKE Shared Secrets

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, 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 and 0.0.0.0 for both values.
3. Enter the IKE Shared Secret and Verify IKE Shared Secret.
4. Click **Done** to apply the configurations.

IKE Policies

1. In the **IKE Policies** section, click **Add** to open the IPsec Add Policy configuration page.
2. Set the Priority to 1 for this configuration to take priority over the Default setting.
3. Set the Encryption type from the drop-down menu.
4. Set the HASH Algorithm to SHA or MD5.
5. Set the Authentication to Pre-Share.
6. Set the Diffie Hellman Group to Group 1 or Group 2.

The IKE policy selections, along with the 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 configuration need to be made on the dialer prior to downloading the dialer onto the local client.

7. Click **Done** to activate the changes.
8. Click **Apply** to apply the changes made before navigating to other pages.

Configure the VPN via the CLI

Authentication Method and Server Addresses

```
vpdn group l2tp
  enable
  ppp authentication {cache-securid|chap|eap|mschap|mschapv2|pap}
  client configuration {dns|wins} <ipaddr1> [<ipaddr2>]
```

Address Pools

```
ip local pool <pool> <start-ipaddr> <end-ipaddr>
```

Source NAT

```
ip access-list session srcnat
  user any any src-nat pool <pool> position 1
```

IKE Shared Secrets

```
crypto isakmp key <key> address <ipaddr> netmask <mask>
```

IKE Policies

```
crypto isakmp policy <priority>
  encryption {3des|aes128|aes192|aes256|des}
  authentication {pre-share|rsa-sig}
  group {1|2}
  hash {md5|sha}
  lifetime <seconds>
```

Example Configurations for Remote Access Clients

This section describes how to configure remote access VPNs for L2TP/IPsec 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 section, L2TP/IPsec requires two levels of authentication: first, IKE SA authentication, and then user-level authentication with a PPP-based authentication protocol. Microsoft clients do not support smart card authentication for the IKE SA. Therefore, the 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 an external RADIUS server using PPP EAP-TLS. In this scenario, 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 following:

- User role for authenticated clients
- RADIUS server and the authentication server group to which the server belongs
- VPN authentication profile which defines the authentication server group and the default role assigned to authenticated clients
- L2TP/IPsec VPN with EAP as the PPP authentication
- IKE policy for preshared key authentication of the SA



Note: 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.

Configuring a VPN for Smart Card Clients via the WebUI

Use the following procedure to configure a L2TP/IPsec VPN for Microsoft smart card clients via the WebUI:

1. Navigate to the **Configuration > Security > Access Control > Policies** page.
2. Click **Add** to add a new policy.
 - a. Enter the name of the policy (for example, authenticated). Default settings for a policy rule permit all traffic from any source to any destination, but you can make a rule more restrictive. You can also configure multiple rules; the first rule in a policy that matches the traffic is applied.
 - b. Click **Add** to add a rule.
 - c. When you are done adding rules, click **Apply**.
 - d. Click the **User Roles** tab. Click **Add** to add a new user role.
 - e. Enter the name of the role (for example, employee).
 - f. Under Firewall Policies, click **Add**. In the Choose from Configured Policies drop-down list, select the policy you previously created. Click **Done**.
 - g. Click **Apply**.

3. Navigate to the **Configuration > Security > Authentication > Servers** page.
 - a. Select **Radius Server** to display the Radius Server List.
 - b. To configure a RADIUS server, enter the name for the server (for example, ias1) and click **Add**.
 - c. Select the name to configure the IP address and key for the server. Select **Mode** to enable the server.
 - d. Click **Apply**.
4. In the Servers list, select **Server Group**.
 - a. Enter the name of the new server group (for example, ias-server) 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 the RADIUS server you just configured from the drop-down menu.
 - e. Click **Add Server**.
 - f. Click **Apply**.
5. Navigate to the **Configuration > Security > Authentication > L3 Authentication** page.
 - a. Select the **default** VPN Authentication Profile.
 - b. From the Default Role drop-down menu, select **employee**.
 - c. Click **Apply**.
 - d. In the **default** VPN Authentication Profile, select **Server Group**.
 - e. Select the server group you just configured from the drop-down menu.
 - f. Click **Apply**.
6. Navigate to the **Configuration > Advanced Services > VPN Services > IPSEC** page.
 - a. Select **Enable L2TP** (this is enabled by default).
 - b. Select **EAP** for Authentication Protocols.
 - c. 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 will be pushed to the VPN client.
 - d. Under Address Pools, click **Add** to open the **Add Address Pool** page.
 - e. Specify the start address, the end address and the pool name.
 - f. Click **Done** to apply the configuration.
 - g. Under IKE Shared Secrets, click **Add** to open the **Add IKE Secret** page.
 - h. To make the IKE key global, specify 0.0.0.0 and 0.0.0.0 for both subnet and subnet mask (these are the default values).
 - i. Enter the IKE Shared Secret and Verify IKE Shared Secret.
 - j. Click **Done** to apply the configurations.
 - k. Under IKE Policies, click **Add** to open the **IPsec Add Policy** configuration page.
 - l. Set the Priority to 1 for this configuration to take priority over the Default setting.
 - m. Set the Authentication to Pre-Share.
 - n. Click **Done** to activate the changes.
 - o. Click **Apply**.

Configuring a VPN for Smart Card Clients via the CLI

Use the following procedure to configure a L2TP/IPsec VPN for Microsoft smart card clients via the CLI:

```
ip access-list session authenticated
  any any any permit position 1
user-role employee
  access-list session authenticated

aaa authentication-server ias1
  host 1.1.1.254
  key 12345678

aaa server-group ias-server
  auth-server ias1

aaa authentication vpn default
  default-role employee
  server-group ias-server

vpdn group l2tp
  enable
  ppp authentication eap
  client dns 101.1.1.245

ip local pool sc-clients 10.1.1.1 10.1.1.250

crypto isakmp key 0987654 address 0.0.0.0 netmask 0.0.0.0

crypto isakmp policy 1
  authentication pre-share
```

VPNs for L2TP/IPsec Clients with 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:

- User role for authenticated clients
- Internal 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
- IKE policy for preshared key authentication of the SA

Configure the L2TP/IPsec VPN via the WebUI

Use the following procedure to configure L2TP/IPsec VPN for username/password clients via the WebUI:

1. Navigate to the **Configuration > Security > Access Control > Policies** page.
2. Click **Add** to add a new policy.
 - a. Enter the name of the policy (for example, authenticated). Default settings for a policy rule permit all traffic from any source to any destination, but you can make a rule more restrictive. You can also configure multiple rules; the first rule in a policy that matches the traffic is applied.
 - b. Click **Add** to add a rule.
 - c. When you are done adding rules, click **Apply**.
 - d. Click the **User Roles** tab. Click **Add** to add a new user role.
 - e. Enter the name of the role (for example, employee).
 - f. Under Firewall Policies, click **Add**. In the Choose from Configured Policies drop-down list, select the policy you previously created. Click **Done**.
 - g. Click **Apply**.
3. Navigate to the **Configuration > Security > Authentication > Servers** page.
 - a. Select **Internal DB** to display entries for the internal database.
 - b. Click **Add User**.
 - c. Enter the username and password.
 - d. Click **Apply**.
4. Navigate to the **Configuration > Security > Authentication > L3 Authentication** page.
 - a. Select **default** VPN Authentication Profile.
 - b. From the Default Role drop-down menu, select employee.
 - c. Click **Apply**.
 - d. Under default VPN Authentication Profile, select Server Group.
 - e. Select the **internal** server group from the drop-down menu.
 - f. Click **Apply**.
5. Navigate to the **Configuration > Advanced Services > VPN Services > IPSEC** page.
 - a. Select **Enable L2TP** (this is enabled by default).
 - b. Select PAP for Authentication Protocols.
 - c. 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 will be pushed to the VPN client.
 - d. Under Address Pools, click **Add** to open the **Add Address Pool** page.
 - e. Specify the start address, the end address and the pool name.
 - f. Click **Done** to apply the configuration.
 - g. Under IKE Shared Secrets, click **Add** to open the Add IKE Secret page.
 - h. To make the IKE key global, specify 0.0.0.0 and 0.0.0.0 for both subnet and subnet mask (these are the default values).
 - i. Enter the IKE Shared Secret and Verify IKE Shared Secret.
 - j. Click **Done** to apply the configurations.
 - k. Under IKE Policies, click **Add** to open the IPSEC Add Policy configuration page.
 - l. Set the Priority to 1 for this configuration to take priority over the Default setting.

- m. Set the Authentication to Pre-Share.
- n. Click **Done** to activate the changes.
- o. Click **Apply**.

Next, you must configure client entries in the internal database.

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 information for the client.
5. Click **Enabled** to activate this entry on creation.
6. Click **Apply** to apply the configuration.

Configure the L2TP/IPsec VPN via the CLI

To configure L2TP/IPsec VPN for username/password clients via the command-line interface, issue the following commands in config mode.

```
ip access-list session authenticated
  any any any permit position 1
user-role employee
  access-list session authenticated

aaa authentication vpn default
  default-role employee
  server-group internal

vpdn group l2tp
  enable
  ppp authentication pap
  client dns 101.1.1.245

ip local pool pw-clients 10.1.1.1 10.1.1.250

crypto isakmp key 0987654 address 0.0.0.0 netmask 0.0.0.0

crypto isakmp policy 1
  authentication pre-share
```

Next, issue the following command in *enable* mode to configure client entries in the internal database:

```
local-userdb add username <name> password <password>
```

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 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.

Configure a VPN with XAuth via the WebUI

Use the following procedures to configure a remote access VPN for Xauth via the WebUI:

Authentication Method and Server Addresses

1. Navigate to the **Configuration > Advanced Services > VPN Services > IPSEC** page.
2. To enable or disable Extended Authentication (XAuth), select or deselect **Enable XAuth** (this is enabled by default).

Disable XAuth if the VPN client is authenticated using a smart card. After successful IKE main mode exchange, the controller extracts the values of the Principal name (SubjectAltname in X.509 certificates) or Common Name fields from the digital certificate in the smart card and authenticates them with the authentication server. The authentication server can be an external RADIUS or LDAP server or the internal database.

3. 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 will be pushed to the VPN client.

Address Pools

This is the pool from which the clients are assigned addresses.

1. Navigate to the **Configuration > Advanced Services > VPN Services > IPSEC** page
2. Under Address Pools, click **Add** to open the **Add Address Pool** page.
3. Specify the start address, the end address and the pool name.
4. Click **Done** to apply the configuration.

Source NAT

Use this option if the IP addresses of clients need to be translated to access the network. To use this option, you must have created a NAT pool by navigating to the **Configuration > IP > NAT Pools** page.

Aggressive Mode

For XAuth clients, the Phase I IKE exchange can be either Main Mode or Aggressive Mode. Aggressive Mode condenses the IKE SA negotiations into three packets (versus six packets for Main Mode). A *group* associates the same set of attributes to multiple clients.

Enter the authentication group name for aggressive mode. Make sure that the group name matches the group name configured in the VPN client software.

Server Certificate

You can specify a single server certificate for VPN clients. The server certificate must be imported into the controller, as described in [Chapter 26, “Management Access” on page 503](#). Select the server certificate from the drop-down list.

CA Certificate for VPN Clients

You can assign one or more trusted CA certificates to VPN clients. The trusted CA certificate must be imported into the controller, as described in [Chapter 26, “Management Access” on page 503](#).

1. Under CA Certificate Assigned for VPN-clients, click **Add**.
2. Select a CA certificate from the drop-down list of CA certificates imported in the controller.
3. Click **Done**.
4. Repeat the above steps to add additional CA certificates.

IKE Shared Secrets

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. Under IKE Shared Secrets, 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 and 0.0.0.0 for both values.
3. Enter the IKE Shared Secret and Verify IKE Shared Secret.
4. Click **Done** to apply the configurations.

IKE Policies

1. Under IKE Policies, click **Add** to open the IPSEC Add Policy configuration page.
2. Set the Priority to 1 for this configuration to take priority over the Default setting.
3. Set the Encryption type from the drop-down menu.
4. Set the HASH Algorithm to SHA or MD5.
5. Set the Authentication to Pre-Share or RSA. If you are using certificate-based IKE, select RSA.
6. Set the Diffie Hellman Group to Group 1 or Group 2.

The IKE policy selections, along with the 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 configuration need to be made on the dialer prior to downloading the dialer onto the local client.

7. Click **Done** to activate the changes.
8. Click **Apply** to apply the changes made before navigating to other pages.

Configure a VPN with XAuth via the CLI

Authentication Method and Server Addresses

```
vpdn group l2tp
  enable
  ppp authentication {cache-securid|chap|mschap|mschapv2|pap}
  client configuration {dns|wins} <ipaddr1> [<ipaddr2>]
{crypto-local isakmp xauth | no crypto-local isakmp xauth}
```

Address Pools

```
ip local pool <pool> <start-ipaddr> <end-ipaddr>
```

Source NAT

```
ip access-list session srcnat
  user any any src-nat pool <pool> position 1
```

Aggressive Mode

```
crypto isakmp groupname <name>
```

Server Certificate

```
crypto-local isakmp server-certificate <name>
```

CA Certificate Assigned for VPN Clients

```
crypto-local isakmp ca-certificate <cacert-name>
```

IKE Shared Secrets

```
crypto isakmp key <key> address <ipaddr> netmask <mask>
```

IKE Policies

```
crypto isakmp policy <priority>
  encryption {3des|aes128|aes192|aes256|des}
  authentication {pre-share|rsa-sig}
  group {1|2}
  hash {md5|sha}
  lifetime <seconds>
```

Configurations 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; in this example, 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:

- User role for authenticated clients
- Entries for Cisco VPN XAuth clients in the controller's internal database



Note: 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.

- VPN authentication default profile which defines the internal authentication server group and the default role assigned to authenticated clients
 - Disable XAuth to disable prompting for the username and password (user credentials are extracted from the smart card)
 - Server certificate to authenticate the controller to clients
 - CA certificate to authenticate VPN clients
- You must install server and CA certificates in the controller, as described in [Chapter 26, “Management Access” on page 503](#).
- IKE policy for RSA (certificate-based) authentication of the SA

Configure a VPN for Cisco XAuth Smart Card Clients

Configure the VPN via the WebUI

The following procedure describes the steps to configure VPN for Cisco Smart Card Clients via the WebUI:

1. Navigate to the **Configuration > Security > Access Control > Policies** page.
2. Click **Add** to add a new policy.
 - a. Enter the name of the policy (for example, authenticated). Default settings for a policy rule permit all traffic from any source to any destination, but you can make a rule more restrictive. You can also configure multiple rules; the first rule in a policy that matches the traffic is applied.
 - b. Click **Add** to add a rule.
 - c. When you are done adding rules, click **Apply**.
 - d. Click the **User Roles** tab. Click **Add** to add a new user role.
 - e. Enter the name of the role (for example, employee).
 - f. Under Firewall Policies, click **Add**. In the Choose from Configured Policies drop-down list, select the policy you previously created. Click **Done**.
 - g. Click **Apply**.

3. Navigate to the **Configuration > Security > Authentication > L3 Authentication** page.
 - a. Select default VPN Authentication Profile.
 - b. From the Default Role drop-down menu, select **employee**.
 - c. Click **Apply**.
 - d. Under default VPN Authentication Profile, select Server Group.
 - e. Select the server group **internal** from the drop-down menu.
 - f. Click **Apply**.
4. Navigate to the **Configuration > Advanced Services > VPN Services > IPSEC** page.
 - a. Select **Enable L2TP** (this is enabled by default).
 - b. Deselect **Enable XAuth** (this is enabled by default).
 - c. 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 will be pushed to the VPN client.
 - d. Under Address Pools, click **Add** to open the **Add Address Pool** page.
 - e. Specify the start address, the end address and the pool name.
 - f. Click **Done** to apply the configuration.
 - g. Select the server certificate the controller will use to authenticate itself to clients.
 - h. Select the CA certificate the controller will use to validate clients. Click **Done**.
 - i. Under IKE Policies, click **Add** to open the IPSEC Add Policy configuration page.
 - j. Set the Priority to 1 for this configuration to take priority over the Default setting.
 - k. Set the Authentication to RSA.
 - l. Click **Done** to activate the changes.
 - m. Click **Apply**.

Next, configure client entries in the internal database:

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 information for the client.
5. Click **Enabled** to activate this entry on creation.
6. Click **Apply** to apply the configuration.

Configure the VPN via the CLI

The following procedure describes the steps to configure VPN for Cisco Smart Card Clients via the CLI:

```
ip access-list session authenticated
  any any any permit position 1
user-role employee
  access-list session authenticated

aaa authentication vpn default
  default-role employee
  server-group internal

no crypto-local isakmp xauth

vpdn group l2tp
```

```

enable
client dns 101.1.1.245

ip local pool sc-clients 10.1.1.1 10.1.1.250

crypto-local isakmp server-certificate ServerCert1
crypto-local isakmp ca-certificate TrustedCA1

crypto isakmp policy 1
  authentication rsa-sig

```

Enter the following command in enable mode to configure client entries in the internal database:

```
local-userdb add username <name> password <password>
```

XAuth Clients Using a Username/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; the user is 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:

- User role for authenticated clients
- Entries for Cisco VPN XAuth clients in the controller's internal database
- VPN authentication profile which defines the internal authentication server group and the default role assigned to authenticated clients
- Enable XAuth to prompt for the username and password
- IKE policy for preshared key authentication of the SA

Configure the VPN via the WebUI

To configure a VPN for Cisco VPN XAuth clients using a username and passwords via the WebUI:

1. Navigate to the **Configuration > Security > Access Control > Policies** page.
2. Click **Add** to add a new policy.
 - a. Enter the name of the policy (for example, authenticated). Default settings for a policy rule permit all traffic from any source to any destination, but you can make a rule more restrictive. You can also configure multiple rules; the first rule in a policy that matches the traffic is applied.
 - b. Click **Add** to add a rule.
 - c. When you are done adding rules, click **Apply**.
 - d. Click the **User Roles** tab. Click **Add** to add a new user role.
 - e. Enter the name of the role (for example, employee).
 - f. Under Firewall Policies, click **Add**. In the Choose from Configured Policies drop-down list, select the policy you previously created. Click **Done**.
 - g. Click **Apply**.
3. Navigate to the **Configuration > Security > Authentication > L3 Authentication** page.
 - a. Select the **default** VPN Authentication Profile.
 - b. From the Default Role drop-down menu, select **employee**.
 - c. Click **Apply**.
 - d. Under the **default** VPN Authentication Profile, select Server Group.
 - e. Select the server group **internal** from the drop-down menu.
 - f. Click **Apply**.

4. Navigate to the **Configuration > Advanced Services > VPN Services > IPSEC** page.
 - a. Select **Enable L2TP** (this is enabled by default).
 - b. Select **Enable XAuth** (this is enabled by default).
 - c. 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 will be pushed to the VPN client.
 - d. Under Address Pools, click **Add** to open the **Add Address Pool** page.
 - e. Specify the start address, the end address and the pool name.
 - f. Click **Done** to apply the configuration.
 - g. Under IKE Shared Secrets, click **Add** to open the **Add IKE Secret** page.
 - h. To make the IKE key global, specify 0.0.0.0 and 0.0.0.0 for both subnet and subnet mask (these are the default values).
 - i. Enter the IKE Shared Secret and Verify IKE Shared Secret.
 - j. Click **Done** to apply the configurations.
 - k. Under IKE Policies, click **Add** to open the **IPSEC Add Policy** configuration page.
 - l. Set the Priority to 1 for this configuration to take priority over the Default setting.
 - m. Set the Authentication to Pre-Share.
 - n. Click **Done** to activate the changes.
 - o. Click **Apply**.

Next, configure client entries in the internal database

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 information for the client.
5. Click **Enabled** to activate this entry on creation.
6. Click **Apply** to apply the configuration.

Configure the VPN via the CLI

To configure a VPN for Cisco VPN XAuth clients using a username and passwords via the CLI:

```
ip access-list session authenticated
  any any any permit position 1
user-role employee
  access-list session authenticated

aaa authentication vpn default
  default-role employee
  server-group internal

crypto-local isakmp xauth

vpdn group l2tp
  enable
  client dns 101.1.1.245

ip local pool pw-clients 10.1.1.1 10.1.1.250

crypto isakmp key 0987654 address 0.0.0.0 netmask 0.0.0.0
```

```
crypto isakmp policy 1
  authentication pre-share
```

Enter the following command in enable mode to configure client entries in the internal database:

```
local-userdb add username <name> password <password>
```

Remote Access VPN 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 as well as 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).

Configuring a VPN with PPTP via the WebUI

1. Navigate to the **Configuration > Advanced Services > VPN Services > PPTP** page.
2. To enable PPTP, select **Enable PPTP**.
3. Select the authentication protocol. The currently-supported method is MSCHAPv2.
4. Configure the primary and secondary DNS servers and primary and secondary WINS Server that will be pushed to the VPN Dialer.
5. Configure the VPN Address Pool.
 - a. Click **Add**. The Add Address Pool page displays.
 - b. Specify the pool name, start address, and end address.
 - c. Click **Done** on completion to apply the configuration.
6. Click **Apply** to apply the changes made before navigating to other pages.

Configuring a VPN with PPTP via the CLI

```
vpdn group pptp
  enable
  client configuration {dns|wins} <ipaddr1> [<ipaddr2>]
  ppp authentication {mschapv2}
pptp ip local pool <pool> <start-ipaddr> <end-ipaddr>
```

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. Or, you can use a VPN concentrator at one site and a controller at the other site .

A Dell controller supports the following IKE SA authentication methods for site-to-site VPNs:

- Preshared key: the same IKE shared secret must be configured on both the local and remote sites.

- Digital certificates: You can configure a server certificate and a CA certificate for each site-to-site VPN IPsec map configuration. For more information about importing server and CA certificates into the controller, see [Chapter 26, “Management Access” on page 503](#).



Note: Certificate-based authentication is only supported for site-to-site VPN between two controllers with static IP addresses.

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. By default, site-to-site VPN uses IKE Main-mode with Pre-Shared-Keys to authenticate the IKE SA. This method uses the IP address of the peer, and therefore will not work for dynamically addressed peers.

To support site-to-site VPN with dynamically addressed devices, you must enable IKE Aggressive-Mode with Authentication based on a Pre-Shared-Key. 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.

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 65 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 the controller A’s interface to the Layer-3 network is located (Interface A in [Figure 65](#))
- The peer gateway, which is the IP address of controller B’s interface to the Layer-3 network (Interface B in the [Figure 65](#))



Note: Configure VPN settings on the controllers at both the local and remote sites.

Configuring site-to-site VPN via the WebUI

Use the following procedure to create a site-to-site VPN via the WebUI:

1. Navigate to the **Configuration > Advanced Services > VPN Services > Site-to-Site** page.
2. Under IPsec Maps, click **Add** to open the Add IPsec Map page.
3. Enter a name for this VPN connection in the **Name** field.
4. Enter the IP address and netmask for the source (the local network connected to the controller) in the **Source Network** and **Source Subnet Mask** fields, respectively. (See controller A in [Figure 65](#))

5. Enter the IP address and netmask for the destination (the remote network to which the local network will communicate) in the **Destination Network** and **Destination Subnet Mask** fields, respectively. (See controller B in [Figure 65](#).)
6. In the **Peer Gateway** field, enter the IP address of the interface on the remote controller that connects to the Layer-3 network. (See Interface B in [Figure 65](#).) 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.
7. The **Security Association Lifetime** parameter defines 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.
8. Select the **VLAN** that contains the interface of the local controller which connects to the Layer-3 network. (See Interface A in [Figure 65](#).)

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).

9. 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 will 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**: Use the 768-bit Diffie Hellman prime modulus group.
 - **group2**: Use the 1024-bit Diffie Hellman prime modulus group.
10. Select **Pre-Connect** to have the VPN connection established even if there is no traffic being sent from the local network. If this is not selected, the VPN connection is only established when traffic is sent from the local network to the remote network.
11. Select **Trusted Tunnel** if traffic between the networks is trusted. If this is not selected, traffic between the networks is untrusted.
12. Select the **Enforce NATT** checkbox to always enforce UDP 4500 for IKE and IPSEC. This option is disabled by default.
13. For 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 a specific initiator only.

14. Select an authentication type. 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 required in IPsec maps for a VPN with a dynamically addressed peer.

-or-

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 [Chapter 26, "Management Access" on page 503](#) for more information.

15. Click **Done** to apply the site-to-site VPN configuration.
16. Click **Apply**.
17. Click the IPsec tab to configure an IKE policy that uses RSA authentication.
 - 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 Encryption type from the drop-down menu.

- d. Set the HASH Algorithm to **SHA** or **MD5**.
- e. Set the Authentication to **PRE-SHARE** if you are using preshared keys. If you are using certificate-based IKE, select **RSA**.
- f. Set the Diffie Hellman Group to **Group 1** or **Group 2**.
- g. The IKE policy selections, along with the 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 the Dell dialer is used, you must configure the dialer prior to downloading the dialer onto the local client.
- h. Click **Done** to activate the changes.
- i. Click **Apply**.

Configuring site-to-site VPN via the CLI

To configure a site-to-site VPN with two static IP controllers via the CLI, issue the following commands:

```
crypto-local ipsec-map <name> <priority>
  src-net <ipaddr> <mask>
  dst-net <ipaddr> <mask>
  peer-ip <ipaddr>
  vlan <id>
  pre-connect enable|disable
  trusted enable
```

For certificates:

```
set ca-certificate <cacert-name>
set server-certificate <cert-name>
```

```
crypto isakmp policy <priority>
  encryption {3des|aes128|aes192|aes256|des}
  authentication rsa-sig
  group {1|2}
  hash {md5|sha}
  lifetime <seconds>
```

For preshared key:

```
crypto-local isakmp key <key> address <ipaddr> netmask <mask>
```

```
crypto isakmp policy <priority>
  encryption {3des|aes128|aes192|aes256|des}
  authentication pre-share
  group {1|2}
  hash {md5|sha}
  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:

```
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:

```
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:

```
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:

```
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:

```
crypto-local isakmp key <key> fqdn-any
```

Dead Peer Detection

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 liveness of an IKE peers.

To configure DPD parameters, issue the following commands via the command-line interface.

```
crypto-local isakmp dpd idle-timeout <idle_seconds> retry-timeout <retry_seconds>
retry-attempts <number>
```

Dell Dialer

For Windows clients, a dialer can be downloaded from the controller to auto-configure tunnel settings on the client.

Configuring the Dialer via the WebUI

Use the following procedure to configure the Dell dialer via the WebUI:

1. Navigate to the **Configuration > Advanced Services > VPN Services > Dialers** page. Click **Add** to add a new dialer or click the **Edit** tab to edit an existing dialer.
2. Enter the Dialer Name that will be used to identify this setting.
3. Configure the dialer to work with PPTP or L2TP by selecting the **Enable PPTP** or the **Enable L2TP** checkbox.
4. Select the authentication protocol. This should match the L2TP protocol list selected if **Enable L2TP** is checked or the PPTP list configured if **Enable PPTP** is checked.
5. For L2TP:
 - Set the IKE Hash Algorithm to SHA or MD5 as in the IKE policy on the **Advanced Services > VPN Services > IPSEC** page.

- If a preshared key is configured for IKE Shared Secrets in the VPN Services > IPSEC page, enter the key.
 - The key you enter in the Dialers page 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.
6. Click **Done** to apply the changes made prior to navigating to another page.

Configuring the Dialer via the CLI

Issue the following commands to configured the Dell dialer via the CLI:

```
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}
```

Captive Portal Dialer

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.

To configure the captive portal dialer via 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**.

To configure the captive portal dialer via the command-line interface, access the CLI in config mode and issue the following commands:

```
user-role <role>
  dialer <name>
```


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 additional 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 WiFi 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 369
- “Configuring Clients” on page 370

Configuring MAC-Based Authentication

Before configuring MAC-based authentication, you must configure:

- The user role that will be assigned as the default role for the MAC-based authenticated clients. (See [Chapter 10, “Roles and Policies”](#) for information on firewall policies to configure roles).
You 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 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 370 for information on configuring the clients on the local database. For information on configuring authentication servers and server groups, see [Chapter 8, “Authentication Servers”](#)

Configuring the MAC Authentication Profile

[Table 67](#) describes the parameters you can configure for MAC-based authentication.

Table 67 MAC Authentication Profile Configuration Parameters

Parameter	Description
Delimiter	Delimiter used in the MAC string: <ul style="list-style-type: none"> • 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 Default: none
Case	The case (upper or lower) used in the MAC string. Default: lower

Table 67 MAC Authentication Profile Configuration Parameters (Continued)

Parameter	Description
Max Authentication failures	Number of times a station can fail to authenticate before it is blacklisted. A value of 0 disables blacklisting. Default: 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 67](#).
6. Click **Apply**.

Using the CLI to configure a MAC authentication profile

```
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 that can be used 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 user name and password for each client.

Note: 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.

Using the WebUI to configure clients in the internal database

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** to apply the configuration.

Note: The configuration does not take effect until you perform this step.

Using the CLI to configure clients in the internal database

Enter the following command in enable mode:

```
local-userdb add username <macaddr> password <macaddr>...
```


ArubaOS supports secure IPsec communications between a controller and campus 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 campus APs. If a local controller is unable to contact the master AP to validate its own certificate, it will not be able to certify its APs, and those APs will not be able to communicate with their local controller until master-local communication has been reestablished. Therefore, it is very important that all local controllers are able to communicate with their master controller when you first enable the control plane security feature.

Some AP model types, such as the W-AP120 Series, have factory-installed digital certificates from Dell PowerConnect. These AP models will use their factory-installed certificates for IPsec, and do not need a self-signed certificate from the controller. Once a campus 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 campus AP will have the same master controller as a common trust anchor. The campus AP whitelist contains a list of all APs connected to the network. You can use this whitelist at any time to add new valid APs to the secure network, or revoke network access to any suspected rogue or unauthorized AP.



Note: The control plane security feature supports campus APs only and is not intended for use with Remote 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.

This chapter describes the following topics:

- [“Control Plane Security Overview” on page 373](#)
- [“Configuring Control Plane Security” on page 374](#)
- [“Whitelists on Master and Local Controllers” on page 380](#)
- [“Environments with Multiple Master Controllers” on page 383](#)
- [“Replacing a Controller on a Multi-Controller Network” on page 386](#)
- [“Troubleshooting Control Plane Security” on page 390](#)

Control Plane Security Overview

Controllers enabled with control plane security will only send certificates to APs that you have identified as valid APs on the network. If you are confident that all campus APs currently on your network are valid APs, you can configure automatic certificate provisioning to send certificates from the controller to each campus AP, or to all campus APs within a specific range of IP addresses. If you want closer control over each AP that gets certified, you can manually add individual campus APs to the secure network by adding each AP's information to the campus AP whitelist.

The default automatic certificate provisioning requires that you manually enter each AP's information into the campus AP whitelist. If you change the automatic certificate provisioning setting to let the controller send

certificates to all APs on the network. This ensures that all valid APs will receive a certificate, but also increases the chance that a rogue or unwanted AP will also be certified. 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 will get a certificate, but any valid AP with an IP address outside the specified address range will not be given a certificate and will not be able to communicate with the controller (except to obtain a certificate). Consider both options carefully before you proceed. If your controller has a publicly accessible interface, you should identify the campus APs on the network by IP address range. This will prevent the controller from sending certificates to external or rogue campus APs that may attempt to access your controller through that publicly accessible interface.

The table below briefly describes the basic tasks to configure the control plane security feature via the campus AP whitelist or automatic certificate provisioning. For complete details on performing each of these individual steps, see [“Configuring Control Plane Security” on page 374](#) and [“Whitelists on Master and Local Controllers” on page 380](#)

Table 68 *Control Plane Security Migration Strategies*

Automatically send Certificates to Campus APs	Manually Certify Campus APs
<p>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.</p>	<p>1. Identify the campus APs that should receive certificates by entering the campus APs' MAC addresses in the campus AP whitelist.</p>
<p>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.</p>	<p>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 “Verify Whitelist Synchronization” on page 390.)</p>
<p>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.</p>	<p>3. Enable the control plane security feature.</p>

Configuring Control Plane Security

Before you enable control plane security for the first time, you must either add all valid APs to the campus AP whitelist *or* enable automatic certificate provisioning. If you do not enable automatic certificate provisioning, only the APs currently approved in the campus AP whitelist will be 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.

Verifying Certificates

Before you begin configuring the control plane security feature on W-600 Series Controller or W-3000 Controller 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 should appear 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=arubanetworks/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 technical support.

```
(host) # show tpm cert-info
Cannot get TPM and Factory Certificate Info.
TPM and/or Factory Certificates might be missing.
```

Configuring Settings via the WebUI

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

Table 69 Control Plane Security Parameters

Parameter	Description
Control Plane Security	Select Enable to enable the control plane security feature. This feature is disabled by default. When control plane security is enabled, any APs on the network that do not have a valid certificate will not be able to communicate with the controller on a clear channel, except to obtain a certificate. NOTE: If you plan on manually adding entries into the AP whitelist, do <i>not</i> enable control plane security until <i>after</i> the completed whitelist has been synchronized to all controllers on the network.
Auto Cert Provisioning	When the control plane security feature is enabled, you can select this checkbox to turn on automatic certificate provisioning. When this feature is enabled, the controller will attempt 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 382. After you have enabled automatic certificate provisioning, you must select either Auto Cert Allow all or Addresses Allowed for Auto Cert .
Auto Cert Allow All	After you enable both control plane security and auto certificate provisioning, select Auto Cert Allow All to allow all associated campus APs to receive automatic certificate provisioning. This parameter is enabled by default.
Addresses Allowed for Auto Cert	If your controller has a publicly accessible interface, you should identify the campus APs by IP address range. This will prevent the controller from sending certificates to external or rogue campus APs that may attempt to access your controller through that interface. After you have enabled both control plane security and auto certificate provisioning, select Addresses Allowed for Auto Cert to send certificates to a group of campus 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 both control plane security and auto certificate provisioning is enabled, all campus APs in the address list will receive automatic certificate provisioning. Remove a range IP addresses from the list of allowed addresses by selecting the IP address range from the list and clicking Delete .

4. Click **Apply** to save your changes.

The master controller will generate its self-signed certificate and will begin 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 66 Control Plane Security Settings

Control Plane Security	<input checked="" type="radio"/> Enable <input type="radio"/> Disabled
Auto Cert Provisioning	<input checked="" type="checkbox"/>
Address Allowed for Auto Cert	<input type="radio"/> All <input checked="" type="radio"/> Specified address range

192.0.2.0 - 192.0.2.20 [Delete] [Add]

Configuring Settings via the CLI

Use the following commands 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 69](#), above.

```
control-plane-security
  auto-cert-allowed-addr <ipaddress-start> <ipaddress-end>
  auto-cert-allow-all
  auto-cert-prov
  cpsec-enable
  no ...
```

Example:

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

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

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

Managing the Campus AP Whitelist

Campus APs appear as valid APs in the campus AP whitelist when you manually enter their information into the whitelist 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 will also be included in the campus AP whitelist, but these APs will appear in an unapproved state.

Use the campus AP whitelist 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 AP whitelist on a controller that uses control plane security, that AP will not be able to communicate with the controller again, except to obtain a new certificate.

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



You can add an AP to the campus AP whitelist 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>AP Installation**.
2. Click the **Campus AP Whitelist** tab.
3. To add a new AP to the whitelist, click **New**.
4. Define the following parameters for each campus AP you want to add to the campus AP whitelist.

Table 70 *Configure Campus AP Whitelist Parameters*

Parameter	Description
AP MAC Address	MAC address of a campus AP that should support secure communications to and from its controller.
Description	(Optional) Use this field to add a brief description of the campus AP.

5. Click **Add** to add the information to the campus AP whitelist.
6. Click **Apply** to save your changes.

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>
```

Viewing Entries in the Campus AP Whitelist

Once you have added an entry in the Campus AP whitelist, that entry will be updated with additional information as the status of the AP changes. To view current information for an AP in the campus AP whitelist via the WebUI:

1. Access the WebUI, and navigate to **Configuration>AP Installation**.
2. Click the **Campus AP Whitelist** tab. The Campus AP whitelist table includes the following information for each AP entry.

Table 71 *View Campus AP Whitelist Parameters*

Parameter	Description
AP MAC Address	MAC address of the campus AP.
Cert Type	The type of certificate used by the AP. <ul style="list-style-type: none"> ● controller-cert: The campus AP is using a certificate signed by the controller. ● factory-cert: the campus AP is using a factory-installed certificate. This option should only be used for AP model types W-AP120 Series.

Table 71 View Campus AP Whitelist Parameters

Parameter	Description
State	<p>The Campus AP Whitelist reports one of the following states for each campus AP:</p> <ul style="list-style-type: none"> ● unapproved-no-cert: AP has no certificate and is not approved. ● unapproved-factory-cert: 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 will not be re-issued a new certificate if automatic certificate provisioning is enabled. ● certified-controller-cert: 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 yet the AP requests to be certified again. Since this is not a normal condition, the AP will 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. <p>NOTE: If an AP is in this state due to connectivity problems, then the AP will recover and will be out of this hold state as soon as connectivity is restored.</p> <ul style="list-style-type: none"> ● certified-hold-controller-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. Since this is not a normal condition, the AP will 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. <p>NOTE: If an AP is in this state due to connectivity problems, then the AP will recover and will be out of this hold state as soon as connectivity is restored.</p>
Description	If defined, a brief description of the campus AP.
Revoked	Shows if the AP's secure status has been revoked.
Revoked Text	An optional, brief statement describing why the AP was revoked.

To view information about the campus AP whitelist via the command-line interface, use the commands described in [Table 72](#).

Table 72 View the Campus AP Whitelist via the CLI

Command	Description
<code>show whitelist-db csec [mac-address <macaddr>]</code>	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.
<code>show whitelist-db csec-status</code>	<p>The command gives aggregate information for the numbers of APs in each of the following categories:</p> <ul style="list-style-type: none"> ● 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 will display 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 are trying to locate, then click the **Search** button. The whitelist will display 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 70](#) 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.
 - **controller-cert:** The campus AP is using a certificate signed by the controller.
 - **factory-cert:** the campus AP is using a factory-installed certificate. This option should only be used for AP model types AP-105 and AP-12x.
 - **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:** AP has been approved state and is ready to receive a certificate.
 - **certified-factory-cert:** 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 will also be allowed to 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 controller-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 will retain 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 will display 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 are trying to locate, then click the **Search** button. The whitelist will display 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 automatic certificate provisioning is enabled for an AP that it is still connected to the network, you will not be able to permanently delete it from the campus AP whitelist; the controller will immediately re-certify the AP and re-create 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 will display 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 are trying to locate, then click the **Search** button. The whitelist will display 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. Any entries in a new controller's campus AP whitelist will be merged into the whitelist for all other master and local controllers as soon as the new controller is added to the hierarchy. If any old or invalid AP entries are added to the campus AP whitelist, all controllers in the hierarchy will begin trusting 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 386](#)

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
```

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 73 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.

Table 73 Control Plane Security Whitelists

Controller Role	Campus AP Whitelist	Master Controller Whitelist	Local Controller Whitelist
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 IP address of the master controller.	The local controller whitelist is empty, and does not appear in the WebUI.

Figure 67 Local Controller Whitelist on a Master Controller

Network > Controller > Control Plane Security

System Settings | Control Plane Security | Cluster Setting | Licenses

Control Plane Security Enable Disabled

Auto Cert Provisioning

Address Allowed for Auto Cert All Specified address range

192.0.2.0 - 192.0.2.20 [Delete] [Add]

AP Whitelist Sync Status

Current sequence number: 3

MAC-Address	IP-Address	Sequence Number	Remote Sequence Number	NULL Update Count	Actions
00:0b:86:01:99:00	10.3.63.2	3	0	0	[Delete]
00:16:CF:AF:3E:E1	172.21.18.170	3	0	0	[Delete]

[Purge]

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 will contain more than one entry. The master controller whitelist rarely needs to be deleted. 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. By default, each controller compares its campus AP whitelist against whitelists on other controllers every two minutes. If a controller detects a difference, it will send 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 **sequence number** column in the local controller or master controller whitelists will have 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 its update acknowledgement may not have yet been received. In the example in [Figure 67](#), 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 [“Verify Whitelist Synchronization” on page 390](#).

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 > Controller**.
2. Select the **Control Plane Security** tab.

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

Table 74 *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: <ul style="list-style-type: none">• 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 currently viewing has received and acknowledged a campus AP whitelist change from the controller in the whitelist. For deployments with both master and local controllers: <ul style="list-style-type: none">• 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. By default, the controller compares its control plane security whitelist against whitelists on other controllers two minutes. If the null update count reaches 5, the controller will send an "empty sync" heartbeat to the remote controller to ensure the sequence numbers on both controllers are the same, then reset 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-ctrl-list [mac-address <mac-address>]
show whitelist-db cpsec-local-ctrl-list [mac-address <mac-address>]
```

Deleting an Entry from the Master or Local Controller Whitelist

There is no 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 will not be physically deleted, as the controller will be waiting for an acknowledgement 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 Controller 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 Controller List For AP Whitelist Sync** section, click the **Delete** button by each controller entry you want to remove.

4. Click **Apply** to save you settings.

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-ctrlr-list del mac-address <mac-address>
```

```
whitelist-db cpsec-local-ctrlr-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 described on [page 382](#).

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 Controller List For AP Whitelist Sync** section, click **Purge**.
- or-
4. To clear the Master Controller Whitelist: In the **Master Controller 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-ctrlr-list purge
```

```
whitelist-db cpsec-local-ctrlr-list purge
```

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 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.




Note: If you add a new backup controller to an existing controller, the backup controller must be added as the **lower priority** controller. If the backup controller is not added 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 will have one master controller as its cluster root, and all other master controllers as cluster members. The master controller operating as the cluster root will create a self-signed certificate, then certify its own local controllers and APs. Next, the cluster root will send a certificate to each cluster member, which in turn certifies their own local controllers and APs. Since all controllers and APs in the cluster will all 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 68 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 for communications between the cluster root and cluster members.

Creating a Cluster Root

To identify a controller as a cluster root via the WebUI:

1. Access the WebUI of the controller you want to become 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 switch 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 save your settings

To create a cluster root via the CLI, access the command-line interface of the controller you want to become the root of the controller cluster, then issue the following command.

```
(host) (config)# cluster-member-ip <ip-address> ipsec <key>
```

The **<ip-address>** parameter is the switch 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.

To identify a controller as a cluster member via 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 have the same value as the key defined for the cluster member in [“Creating a Cluster Root” on page 384](#).
6. Click **Add**.
7. Click **Apply** to save your settings.

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

```
(host) (config)# cluster-root-ip <ip-address> ipsec <key>
```

In this command the **<ip-address>** parameter is the IP address of the root master controller in the cluster, and the **<key>** parameter in each command is the IPsec pre-shared key for communication between the member and the cluster root. This parameter must 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 75](#).

Table 75 CLI Commands to Display Cluster Settings

Command	Description
<code>show cluster-switches</code>	When you issue this command from the cluster <i>root</i> , the output of this command displays the IP address of the VLAN used by the cluster member to connect to the cluster root. If you issue this command from a cluster <i>member</i> , the output of this command displays the IP address of the VLAN used by the cluster root to connect to the cluster member.
<code>show cluster-config</code>	When you issue this command from the cluster <i>root</i> , 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 <i>member</i> , 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.

Note: The following sections describe the steps to replace an existing controller. To add a new local controller to a network, or permanently remove a local controller without replacing it, see [“Viewing and Managing the Master or Local controller Whitelists” on page 382](#).

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 the campus AP whitelist has been purged, you must inform the master controller that the local controller will no longer be available. .

Note: 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-ctrl-list del mac-address <local-controller-mac>`

-or -

Access the maser 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. Access the command-line interface on the new 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**.
5. Now, connect the new local controller to the network. It is very important that the local controller is able to contact the master controller the first time it is connected to the network, because the local controller will try to get its control plane security certificate certified by the master controller the first time the local controller contacts its master.
6. Once the local controller has a valid control plane security certificate and configuration, the local controller will receive the campus AP whitelist from the master controller and will start certifying approved APs.
7. APs associated with the new local controller will reboot and create 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 will generate 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 that the local controllers get their certificate from the new master. Each local controller will begin using a new certificate signed by the master controller.
5. APs will no longer be able to securely communicate with the controller using their current key, and must receive 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 will send the approved APs new certificates. The APs will 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 at running. This will ensure that all certificates, keys and whitelist entries are synchronized to the backup controller. Since the AP whitelist may change periodically, the network administrator should regularly synchronize these settings to the backup controller. For details, see [“Configuring Networks with a Backup Master Controller” on page 383](#)

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 will be 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 386](#)

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 will 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 will need to reboot and get recertified before they can connect to a different controller. If the cluster member you are removing has local controllers, the local controllers will also reboot so they can update themselves 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 will get an updated certificate.
4. If the cluster member you removed has any associated local controllers, reboot those local controllers so they can get 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-ctrl-list del mac-address <cluster-master-mac>`.



Note: 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 will still have 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 385](#). 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 to allow them to get 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 will obtain a new certificate signed by the cluster member, and will 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 at running. This will ensure that all certificates, keys and whitelist entries are synchronized to the backup controller. Since the AP whitelist may change periodically, the network administrator should regularly synchronize these settings to the backup controller. For details, see [“Configuring Networks with a Backup Master Controller” on page 383](#).

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 will be 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 will create its own self-signed certificate. You will 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 will begin using a new certificate signed by the cluster root.
6. Reboot every local controller. Each local controller will begin using a new certificate signed by the cluster member.
7. Because the cluster root is new, it will 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 will 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

Dell recommends using 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 will ensure that all certificates, keys and whitelist entries are synchronized to the backup controller. Since the AP whitelist may change periodically, the network administrator should regularly synchronize these settings to the backup controller. For details, see [“Configuring Networks with a Backup Master Controller” on page 383](#).

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 will be 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

Troubleshooting Control Plane Security

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-controller-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 yet the AP requests to be certified again. Since this is not a normal condition, the AP will 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 will recover and will be out of this hold state as soon as connectivity is restored.
- **certified-hold-controller-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. Since this is not a normal condition, the AP will 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 will recover and will be out of this hold state as soon as connectivity is restored.

Disabling Control Plane Security

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

If you disable control plane security on a *master* controller, APs directly connected to the master controller will reboot then reconnect to the master controller over a clear channel. However, its local controllers will 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 will also reboot and reconnect over a secure channel.

Verify 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 69 Sequence numbers on Master and Local Controllers

Master Controller

Control Plane Security Enable Disabled
 Auto Cert Provisioning
 Address Allowed for Auto Cert All Specified address range

172.1.1.1 - 172.1.1.5
 192.1.1.1 - 192.1.1.5
 198.1.1.1 - 198.1.1.5
 198.1.1.6 - 198.1.1.10
 199.1.1.1 - 199.1.1.5
 199.1.1.6 - 199.1.1.10
 22.22.22.12 - 22.22.22.24
 22.22.22.22 - 22.22.22.24
 122.22.22.12 - 122.22.22.24

AP Whitelist Sync Status
Current sequence number: 92

Local Controller List For AP Whitelist Sync							
MAC-Address	IP-Address	Sequence Number	Remote Sequence Number	NULL Update Count	Actions		
00:0b:86:61:8f:70	10.4.21.33	92	175	1	<input type="button" value="Delete"/>		
00:0b:86:f0:10:16	10.4.21.200	92	148	1	<input type="button" value="Delete"/>		

Local Controller

System Settings | Control Plane Security | Licenses

Control Plane Security Enable Disabled
 Auto Cert Provisioning
 Address Allowed for Auto Cert All Specified address range

AP Whitelist Sync Status
Current sequence number: 148

Master Controller List For AP Whitelist Sync							
MAC-Address	IP-Address	Sequence Number	Remote Sequence Number	NULL Update Count	Actions		
00:0b:86:61:10:4c	10.4.21.34	148	92	4	<input type="button" value="Delete"/>		

Supported APs

The control plane security feature is supported on W-AP120 Series and Ortronics APs. APs that do not support control plane security will not be able to connect to a controller enabled with this feature.

Rogue APs

If you enable auto certificate provisioning enabled with the **Auto Cert Allow All** option, any AP that appears on the network will receive 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.

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 393
- “[Configuring Local Controllers](#)” on page 395

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

Preshared Key for Inter-Controller Communication

A preshared key (PSK) is used 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.



Note: 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. See “[Best Security Practices for the Preshared Key](#)” on page 394. 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. See “[Best Security Practices for the Preshared Key](#)” on page 394. 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.

Best Security Practices for the Preshared Key



Note: 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.

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.

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 ~'@\$%^&*()_+ = \ | / . [] { }

Configuring the Preshared Key

The following sections describe how to configure a PSK using the WebUI or CLI.

Using the WebUI to configure the 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 the Master Controller PSK

Use the procedure below to configure 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.
 - (Optional) In the **FQDN** field, enter a fully qualified domain name used in IKE.
 - (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.
3. Click **Apply**.

Using the CLI to configure the 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:



Note: 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 Local Controllers

A single master controller configuration can be one controller or a master redundant configuration with one master controller and the VRRP redundant backup controller. 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.

Configuring the Local Controller

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 *ArubaOS Quick Start Guide* and are referred to throughout this chapter 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.



Note: 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 preshared key (PSK) that is used to authenticate communications between controllers.



Note: 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 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>
```


A *mobility domain* is a group of Dell controllers among which a wireless user 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.

This chapter describes the following topics:

- “Dell Mobility Architecture” on page 399
- “Configuring Mobility Domains” on page 400
- “Tracking Mobile Users” on page 404
- “Advanced Mobility Functions” on page 406
- “Mobility Multicast” on page 409

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, an 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 where the client appears for the first time when it joins 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 70 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 70 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 70 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 Dell recommends 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 402](#).

Table 76 Basic tasks to configure 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)	<ul style="list-style-type: none"> 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 IP mobility is enabled 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.

The home agent table (HAT) maps a user VLAN IP subnet to potential home agent addresses. 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.)

Dell recommends 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.



Note: 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 home network.

Configure the HAT with a list of every subnetwork, mask, VLAN ID, VRRP IP, and home agent IP address in the mobility domain. Include an entry for every home agent and user VLAN to which an IP subnetwork maps. If there is more than one controller in the mobility domain providing service for the same user VLAN, you must configure an entry for the VLAN for each controller. Dell recommends using the same VRRP IP used by the AP.

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>
    nat <subnetwork> <netmask> <vlan-id> <home-agent-address>
```

The VLAN ID must be the VLAN number on the home agent controller.

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, including surveying 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 are using 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

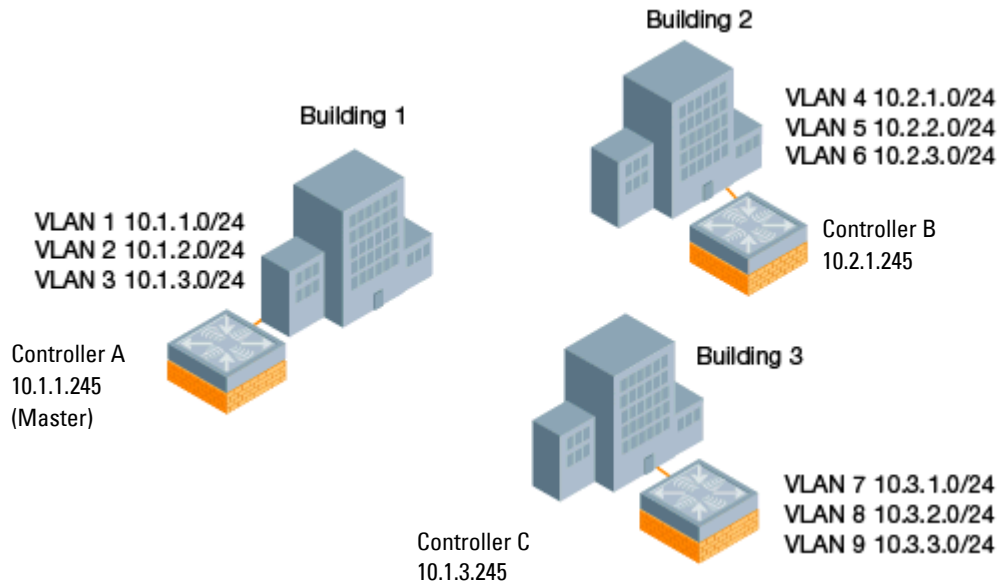
```
ip mobile active-domain <name>
```

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 71](#)) configures a network in a campus with three buildings. An 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 (controller A in this example). On the local controllers (controllers B and C), you only need to enable mobility.

Figure 71 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. Select the **Enable IP Mobility** checkbox.
2. Select the “default” domain in the Mobility Domain list.
3. Click **Add** under the Subnet column. Enter the subnetwork, mask, VLAN ID, and home agent IP address for the first entry shown below and click **Add**. Repeat this step for each HAT entry.

Table 77 Example entries

Subnetwork	Mask	VLAN ID	Home Agent Address or VRIP
10.1.1.0	255.255.255.0	1	10.1.1.245
10.1.1.0	255.255.255.0	1	10.2.1.245
10.1.2.0	255.255.255.0	2	10.1.1.245
10.1.3.0	255.255.255.0	3	10.1.1.245
10.2.1.0	255.255.255.0	4	10.2.1.245
10.2.2.0	255.255.255.0	5	10.2.1.245
10.2.3.0	255.255.255.0	6	10.2.1.245
10.3.1.0	255.255.255.0	7	10.3.1.245
10.3.2.0	255.255.255.0	8	10.3.1.245
10.3.3.0	255.255.255.0	9	10.3.1.245

4. Click **Apply**.

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):

```
ip mobile domain default
  hat 10.1.1.0 255.255.255.0 1 10.1.1.245
  hat 10.1.1.0 255.255.255.0 1 10.2.1.245
  hat 10.1.2.0 255.255.255.0 2 10.1.1.245
  hat 10.1.3.0 255.255.255.0 3 10.1.1.245
  hat 10.2.1.0 255.255.255.0 4 10.2.1.245
  hat 10.2.2.0 255.255.255.0 5 10.2.1.245
  hat 10.2.3.0 255.255.255.0 6 10.2.1.245
  hat 10.3.1.0 255.255.255.0 7 10.3.1.245
  hat 10.3.2.0 255.255.255.0 8 10.3.1.245
  hat 10.3.3.0 255.255.255.0 9 10.3.1.245
router mobile
```

On controllers B and C:

```
router mobile
```

Tracking Mobile Users

This section describes the ways in which 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 user name, role, and authentication can be different on the home agent and foreign agent, as explained by the following: 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 re-authentication 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

```
show ip mobile host
```


Roaming status can be one of the following:

Table 78 *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.

You can view the roaming status of users on any controller in the mobility domain:

Viewing user roaming status using the CLI

```
show user
```

Roaming status can be one of the following:

Table 79 *User Roaming status*

Status Type	Description
Associated	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.

You can use the following CLI command to view the home agent, foreign agent, and roaming status for a specific mobile client.

Viewing specific client information using the CLI

```
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

```
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, etc. This delays HA discovery and eventually resulting in loss of downstream traffic if any meant for the mobile client.

With HA discovery on association, a controller can perform a HA discovery as soon as the client is associated. This feature can be enabled using the `ha-disc-onassoc` parameter in the `wlan virtual <ap-profile>` command. By default, this feature is disabled. You can enable this on virtual APs with devices in power-save mode and requiring mobility. This option will also poll for all potential HAs.

Setting up mobility association using CLI

```
wlan virtual-ap default ha-disc-onassoc
```

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.

Configuring advanced mobility functions using 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 80](#) describes the parameters you can configure on the **Global Parameters** tab.

Table 80 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". The range of allowed values is 10-65534 seconds. The default setting is 180 seconds.
Max. Visitors Allowed	Set a maximum allowed number of active visitors. The range of allowed values for this option is 0-5000 visitors. The default setting is 5000 visitors.

Table 80 IP Mobility Configuration Parameters

Parameter	Description
Registration Requests Retransmits	Maximum number of times the foreign agent attempts mobile IP registration message exchanges before giving up. The range of allowed values for this option is 0-5 attempts. The default setting is 3 attempts.
Registration Requests Interval	Retransmission interval, in milliseconds. The range of allowed values for this option is 100-10000 milliseconds, inclusive. The default setting is 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. The range of allowed values is 0-5000 seconds. The default setting is 5000 seconds.
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. The range of allowed values is 0-300 seconds. The default setting is 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 that can trigger mobility. The downside to this option is lowered security; an association is all it takes to trigger mobility, however, this is irrelevant unless layer-2 security is enforced.
Stand Alone AP Support	Enables support for third party or standalone APs. When this is enabled, broadcast packets are not used to trigger mobility and packets from untrusted interfaces are accepted. If mobility is enabled, you must also enable standalone AP for the client to connect to the controller's untrusted port. If the controller learns wired users via the following methods, enable standalone AP: <ul style="list-style-type: none"> • Third party AP connected to the controller through the untrusted port. • Clients connected to ENET1 on the W-AP92. • Wired user connected directly to the controller's untrusted port. NOTE: When IP mobility is enabled, you must also enable the Stand Alone AP Support option so that a Mux server can perform properly and display all wired users who are connected to a Mux port.
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.
Blocking DHCP Release from stations	Determines whether DHCP release packets generated from the client should be dropped or forwarded to the DHCP server. Blocking the packets prevents the DHCP server from assigning the same IP address to another client until the lease has expired.
Re-Homing for Voice Capable Client	Allows on-hook phones to be assigned a new home agent. This is to load balance voice client home agents across controllers in a mobility domain. This parameter requires that you install the Policy Enforcement Firewall Next Generation (PEFNG) license in the controller, and is disabled by default.
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. The allowed range of values is 1-65535 events, and the default value is 25 events.
Station Trail Timeout	Specifies the maximum interval, in seconds, an inactive mobility trail is held. The allowed range of values is 120-86400 seconds, and the default value is 3600 seconds.

Table 80 IP Mobility Configuration Parameters

Parameter	Description
Station Trail Max. Entries	Specifies the maximum number of entries (client moves) stored in the user mobility trail. The allowed range of values is 1-100 entries, and the default value is 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. Note that in many case a station state is deleted without waiting for the stale timeout; user delete from management, foreign agent to foreign agent handoff, etc. (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.
Proxy DHCP	
Max. BOOTP Messages per Transaction	Maximum number of BOOTP packets that are allowed to be handled during one DHCP session. The allowed range of values for tis parameter is 0-65534 packets. The default value is 25.
Max. Time Allowed per DHCP Transaction	Maximum time allowed for a proxy DHCP session to complete. The allowed range of values for this parameter is 1-600 seconds. The default value is 60 seconds.
Time to hold DHCP state after transaction completion	Hold time, in seconds, on proxy DHCP state after completion of DHCP transaction (DHCP ACK) was forwarded to the client. This option ensures that late BOOTP replies reach the station and that a retransmitted BOOTP request does not trigger a new proxy DHCP session. The allowed range of values for this parameter is 1-600 seconds. The default value is 5 seconds.
Revocation	
Retransmits	Maximum number of times the home agent or foreign agent attempts mobile IP registration/ revocation message exchanges before giving up. The allowed range of values for this parameter is 0-5 retransmissions. The default value is 3 retransmissions.
Interval	Retransmission interval, in milliseconds. The allowed range of values for this parameter is 100-10000 milliseconds. The default value is 1000 milliseconds.

4. Click **Apply** after setting the parameter.

Configuring mobility functions using CLI

```
ip mobile foreign-agent {lifetime <seconds> | max-visitors <number> |
  registrations {interval <msecs> | retransmits <number>}}
```

```
ip mobile home-agent {max-bindings <number>|replay <seconds>}
```

```
ip mobile proxy auth-sta-roam-only | block-dhcp-release | dhcp {max-requests
  <number>|transaction-hold <seconds>|transaction-timeout <seconds>}| event-threshold
  <number> | log-trail | no-service-timeout <seconds> | on-association |re-home |
  stale-timeout <seconds> | stand-alone-AP | trail-length <number> |trail-timeout
  <seconds>
```

```
ip mobile revocation {interval <msec>|retransmits <number>}
```

```
ip mobile trail {host IP address | 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.

Proxy DHCP

When a mobile client first associates with a controller, it sends a DHCP discover request with no requested IP. The controller allows DHCP packets for the client onto the configured VLAN where, presumably, it will receive an IP address. The incoming VLAN becomes the client's home VLAN.

If a mobile client moves to another AP on the same controller that places the client on a different VLAN than its initial (home) VLAN, the *proxy DHCP module* redirects packets from the client's current/visited VLAN to the home VLAN. The proxy DHCP module also redirects DHCP packets for the client from the home VLAN to the visited VLAN.

If the mobile client moves to another controller, the proxy DHCP module attempts to discover if the client has an ongoing session on a different controller. When a remote controller is identified, all DHCP packets from the client are sent to the home agent where they are replayed on the home VLAN. The proxy DHCP module also redirects DHCP packets for the client from the home VLAN to the visited network. In either situation, operations of the proxy DHCP module do not replace DHCP relay functions which can still operate on the client's home VLAN, either in the controller or in another device.

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.

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. ArubaOS supports mobility multicast enhancements that provide uninterrupted streaming of multicast traffic, regardless of a client's location.

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 Dell controller running IGMP proxy as a host; a client attached to the controller would see the controller as

router. When Proxy IGMP is enabled, all multicast clients not associated with the controller are hidden from the upstream multicast device or router.




Note: 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 of the 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 ip address or by upstream port so the controller can dynamically update the upstream multicast router information.

Inter-controller Mobility

When a client moves from one controller to another, multicast traffic migrates as follows:



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 IPIP tunnel. The remote controller will receive downstream multicast traffic and send it to the mobile client.

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 it 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 he 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 check to see if any of their other clients require group1 traffic. If none of their other clients are interested in 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 group membership.

Configuring Mobility Multicast Using the WebUI

To configure the mobility multicast feature using the controller WebUI:

1. Navigate to the **Configuration > Network > IP** window. Click the **Edit** button by the VLAN interface for which you want to configure mobility multicast. The **Edit VLAN** window opens.
2. Select **Enable IGMP** to enable the router to discover the presence of multicast listeners on directly-attached links.
3. Select **Snooping** to save bandwidth and limit the sending of multicast frames to only those nodes that need to receive them.
4. 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.
5. Click **Apply** to apply your changes.

(Optional) Repeat steps 1-6 above to configure mobility multicast for another VLAN interface.

Configuring Mobility Multicast Using 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 81 *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>	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>.
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
```


The underlying mechanism for the Dell redundancy solution is the Virtual Router Redundancy Protocol (VRRP). 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
- A pair of controllers acting as a redundant pair of master controllers in a hot-standby mode

VRRP eliminates a single point of failure by providing an election mechanism, among the controllers, to elect a VRRP “master” controller. The master controller election is:

- If VRRP preemption is disabled (the default setting) and all controllers share the same priority, the first controller that comes up becomes the master.

or

- If VRRP preemption is enabled and all controllers share the same priority, the controller with the highest IP address becomes the master.

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.

Redundancy Parameters

Depending on your redundancy solution, you configure the VRRP parameters listed in [Table 82](#) on your master and local controllers.

Table 82 VRRP Parameters

Parameter	Description
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	This is the interval, in seconds, between successive VRRP advertisements sent by the current <i>master</i> . The default interval time is recommended. Default: 1 second
Authentication Password	This is an optional password, of up to eight characters, that can be used to 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 Address	This is the virtual IP address that will be owned by the elected VRRP <i>master</i> .
Enable Router Pre-emption	Selecting this option means that a controller can take over the role of <i>master</i> if it detects a lower priority controller currently acting as <i>master</i> .
Priority	Priority level of the VRRP instance for the controller. This value is used in the election mechanism for the <i>master</i> .

Table 82 VRRP Parameters (Continued)

Parameter	Description
Tracking	<p>Configures a tracking mechanism that modifies a specified <i>value</i> 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.</p> <p>Tracking can be based on one of the following:</p> <ul style="list-style-type: none"> • Master Up Time: how long the controller has been the master. The value of <i>duration</i> 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. <p>Tracking can also be based on the interface states of the controller:</p> <ul style="list-style-type: none"> • VLAN and Interface: prevents asymmetric routing by tracking multiple VRRP instances. The priority of the VRRP interface determined by the <i>sub</i> 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. <p>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.</p>
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 will run.

Configuring the Local Controller for Redundancy

In an Dell network, the APs are controlled by a controller. The APs tunnel all data to the controller which processes the data, including encryption/decryption, bridging/forwarding, etc.

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.

Note: 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, A800 to A800 or higher), and both controllers should be running the same version of ArubaOS.

The APs are then configured to connect to the “virtual-IP” configured for the VRRP instance.

Collect the following information needed to configure local controller redundancy:

- **VLAN ID** on the two local controllers that are on the same Layer-2 network and is used to configure VRRP.
- **Virtual IP address** to be used for the VRRP instance.

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. Enter the IP Address for the virtual router. Select the VLAN on which VRRP will run. Set the Admin State to Up.

4. Click **Done** to apply the configuration and add the VRRP instance.

In the CLI

```
vrrp <id>  
  ip address <ipaddr>  
  vlan <vlan>  
  no shutdown
```

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 [“Configuring APs” on page 396](#).



Note: 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:

```
ap system-profile <profile>  
  lms-ip <ipaddr>  
  
ap-group <group>  
  ap-system-profile <profile>  
  
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, RF configuration to ease the configuration and maintenance of a wireless network. It also maintains a database related to the wireless network that is used to make any 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 and will be used to configure VRRP.
 - **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 82](#)). 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*”.

```
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.

```
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 83 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.



Note: All the APs and local controllers in the network should be configured with the virtual IP address as the master IP address. The master IP address can be configured for local controllers during the Initial Setup (refer to the ArubaOS Quick Start Guide). You can also use the following commands to change the master IP of the local controller. 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. In addition, you can also synchronize RF Plan data between the pair of controllers. You can either manually or automatically synchronize the databases.



Note: When synchronizing the databases, Dell recommends that you also synchronize RF plan data.

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. Dell recommends a minimum value of 20 minutes.
 - c. By default, RF Plan data is also synchronized. Dell recommends that you always enable this option.
3. Click **Apply**.

In the CLI

Use the following commands to configure database synchronization.

Command	Description
<code>database synchronize</code>	This enable mode command manually synchronizes the databases and takes effect immediately.
<code>database synchronize rf-plan-data</code>	This config mode command includes RF plan data when synchronizing databases. This data is included by default.
<code>database synchronize period <minutes></code>	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:

```
show database synchronize
```

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.



Note: This solution requires that the master controller have Layer-2 connectivity to all the local controllers.

Figure 72 Redundant Topology: Master-Local Redundancy



The network in [Figure 72](#), 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 82](#)). 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



Note: The master controller is configured for a number of VRRP instances (equal to the number of local controllers the master is backing up).

The following example configuration of the master controller in such a topology for one of the VLANs (in this case VLAN 22).

```
vrrp 22
  vlan 22
  ip address 10.200.22.254
  priority 100
  preempt
  authentication password
  description Master-acting-as-backup-to-local
  tracking master-up-time 30 add 20
  no shutdown
```

The following example configuration on the corresponding local controller.

```
vrrp 22
  vlan 22
  ip address 10.200.22.254
  priority 110
  preempt
  authentication password
  description local-backed-by-master
  no shutdown
```

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.

As an example, the administrator can configure 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.



Note: You configure APs 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.

Dell's 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.

Migration and Interoperability

Dell's RSTP implementation interoperates with both 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

Rapid Convergence

Since RSTP is backward compatible with STP, it is possible to configure bridges RSTP (and STP) 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 84](#) compares the port states between STP and RSTP.

Table 84 *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 85](#)).

Table 85 *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, by default, point-to-point links.

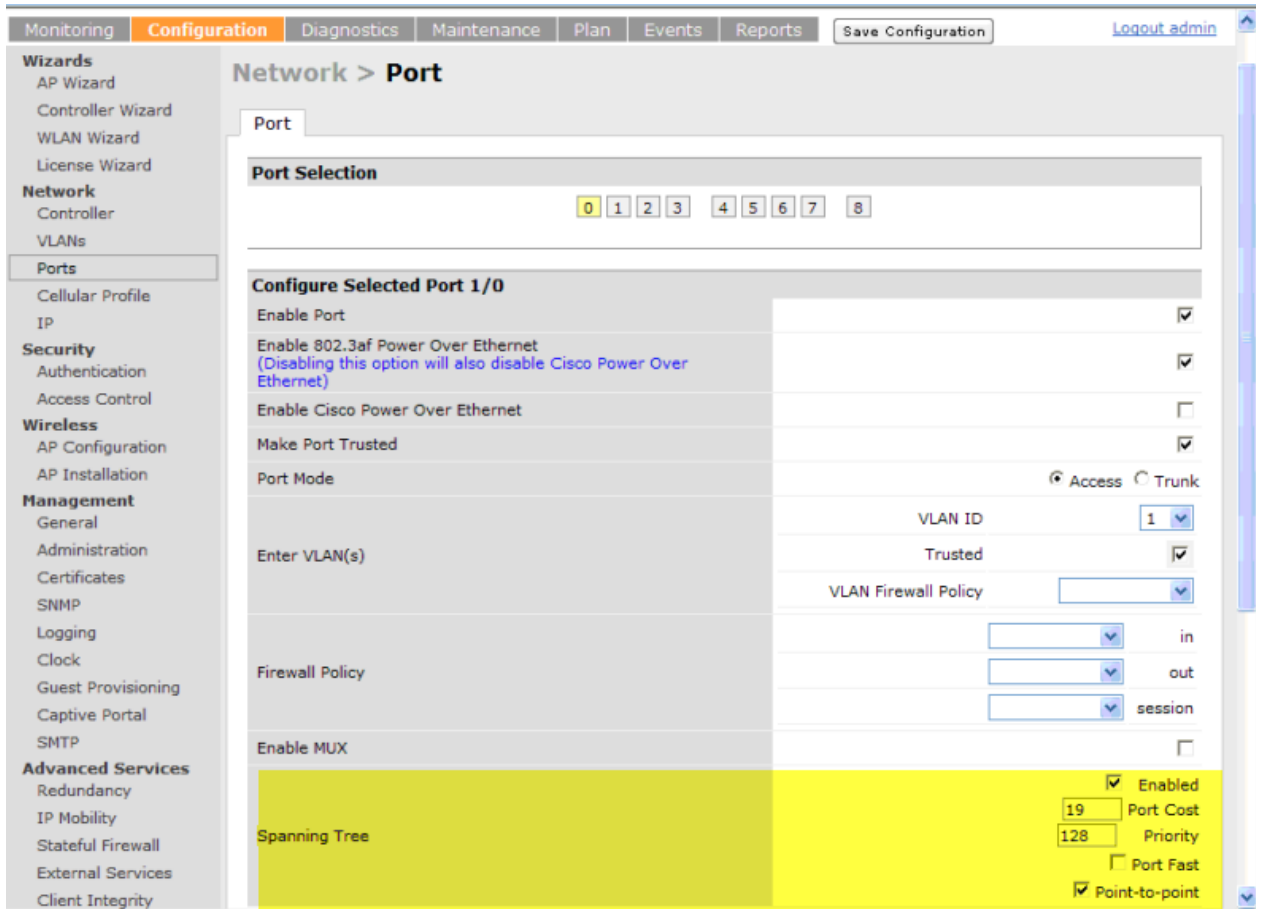
Configuring RSTP

Use either the command line reference or the WebUI to configuring RSTP.

In the WebUI

The RSTP port interface is designated as point-to-point, by default, in the existing port configuration screen ([Figure 73](#)).

Figure 73 Configuring RSTP



Since RSTP is enabled by default, the default values appear in the WebUI. Table 86 list the RSTP defaults and ranges (when applicable) in the configuration interface mode (config-if).

Table 86 RSTP Default Values

Feature	Default Value/Range
Port Cost	The RSTP interface path cost. Range: 1 - 65536 Default: Based on Interface type: Fast Ethernet 10Mbps—100 Fast Ethernet 100Mbps—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	Enabled—Set the interface as a point-to-point link

In the CLI

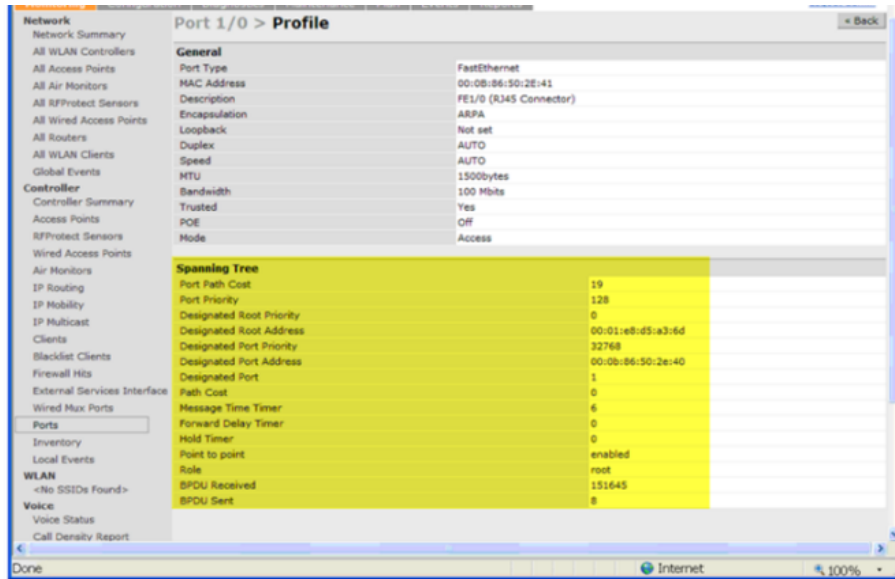
Change the default configurations via the command line.

```
(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 74](#)).

Figure 74 Monitoring RSTP



Troubleshooting

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, etc.); 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, Tx counter will keep incrementing while Rx counter will remain the same. It is quite 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)#
```


The Dell PowerConnect W-600 Series Controller is designed for compact, cost-effective "all-in-one" networking solutions. The W-600 Series Controller includes a firewall, wireless LAN controller, 9-port (8-port for the W-650 and W-651) Ethernet switch with PoE+, IP router, site-to-site VPN edge device, file server, and print server. Additionally, the W-651 controller includes an integrated single radio dual-band (802.11 a/n or 802.11 b/g/n) wireless internal Access Point (AP).

The W-600 Series Controller is an enterprise-class, wireless LAN controller that connects, controls, and integrates wireless APs and Air Monitors (AMs) into a wired LAN system. [Table 87](#) list some of the hardware features by the numbers.

Table 87 W-600 Series Controller by the Numbers

Controller	USB Ports	Maximum External APs	Internal AP	Remote APs
W-620	1	8	None	32
W-650	4	16	None	64
W-651	4	16	1	64

The sections in this chapter are:

- [“Important Points to Remember” on page 427](#)
- [“Internal Access Point \(AP\)” on page 428](#)
- [“USB Cellular Modems” on page 428](#)
- [“Configuring a Supported USB Modem” on page 432](#)
- [“Configuring a New USB Modem” on page 433](#)
- [“NAS \(Network-Attached Storage\)” on page 437](#)
- [“Print Server” on page 442](#)
- [“Sample Topology and Configuration” on page 444](#)

Important Points to Remember

- Only FAT16, FAT32, ext2 and ext3 partitions are supported.
- For shared folders in an ext2/ext3 partition, the owner of the folder must be “nobody”. Otherwise clients will not be able to access the shared folder.
- Unsupported partitions may exist on the NAS device; only supported partitions are mounted.
- User authentication for file access is not supported. The same permissions are applicable to all users.
- Sharing disks that contain errors may cause unpredictable behavior. Scan the disk for errors before mounting the disks to a W-600 Series Controller.
- Un-mount all partitions before disconnecting the disk from the controller.
- Detection of devices connected to an external USB hub may be unpredictable.
- A USB hard disk connected to the controller via an USB ExpressCard adapter is not supported

Internal Access Point (AP)

The W-651 controller includes an internal AP. The internal AP is provisioned in the same way as any other external AP. The provisioning data is stored in the NVRAM. The internal AP identifies itself to a Master controller as the W-651. The internal AP can operate as an AP, Mesh Portal, or an Air Monitor. However, the W-651 internal AP can not operate as a remote AP, a mesh point, or an RF Protect sensor.

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) modems. The 3G HSPA is provided by AT&T in the United States and numerous other 3G providers worldwide.

Functional Description

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.

Mode-Switching

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.

USB Modems Commands

To support the USB cellular modems on the W-600 Series Controller, cellular specific commands are available at the command line (see [Figure 75](#) and [Figure 76](#)). For detailed information on these commands, refer to the Command Line Reference Guide.

Figure 75 Cellular Profile Commands

```
(host) (config) # cellular profile profile_name
(host) (config-cellular profile_name)# ?
dialer          Dialer group settings
driver          Cellular modem driver
import          Import USB device parameters
modeswitch      USB device modeswitch settings
no              Delete Command
priority        Override default priority
serial          USB device serial
tty             Modem TTY port
user            User name authentication
vendor          USB Vendor ID

(host) (config-cellular profile_name)#
```

Figure 76 list the Uplink commands.

Figure 76 Uplink Commands

```
(host) (config) # uplink ?
cellular        Cellular uplink configuration
disable         Disable uplink manager
enable          Enable uplink manager
wired           Wired uplink configuration

(host) (config) # uplink
```

You can view connected USB cellular devices via the **Controller > Universal Serial Bus > USB Devices** in the Web UI (see Figure 77). Navigating to this page is the equivalent of executing the `show usb` command at the command prompt.

Figure 77 Connected Cellular Devices

The screenshot shows the Web UI navigation menu on the left with 'Monitoring' selected. The main content area is titled 'Controller > Universal Serial Bus > USB Devices'. Below the title are tabs for 'USB Devices', 'USB Details', and 'USB Storage'. A timestamp indicates the last update at 13/7/2009 15:46:33. A table titled 'USB Device Table' contains the following data:

Address	Product	Vendor	ProdID	Serial	Type	Profile	Stat
3	Flash Voyager	1b1c	0ab1	a3c60203be7320	Storage		

Uplink Manager

Access the Uplink Manager feature from the WebUI Configuration tab. Navigate to this feature via **Uplink > Uplink Manager** (Figure 78).

Figure 78 WebUI Uplink Manager

Monitoring Configuration Diagnostics Maintenance Plan Events Reports [Logout](#)

Network

- Network Summary
- All WLAN Controllers
- All Access Points
- All Mesh Nodes
- All Air Monitors
- All Wired Access Points
- All Routers
- All WLAN Clients
- Global Events

Controller

- Controller Summary
- Access Points
- Mesh Nodes
- Wired Access Points
- Air Monitors
- IP Routing
- IP Mobility
- IP Multicast
- Clients
- Blacklist Clients
- Wired Mux Ports
- Ports
- Uplink**
- Universal Serial Bus
- Inventory

Controller > Uplink > Uplink Management

Last update at: 13/7/2009 15:52:43

Uplink Management Table

	Id	Uplink Type	Properties	Priority	State	Status
	1	Wired	vlan 1	200	Connected	* Active *

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 79](#)) the bottom modify portion is revealed. All changes are entered into the buffer until the Apply button is executed.

Figure 79 Cellular Profile from the WebUI

Network > Uplink > Cellular Profile

Uplink Manager Cellular Profile Dialer Group

Cellular Profile Table										
Name	Vend	Prod	Serial	Dialer	Tty	Driver	Priority	Modeswitch	Actions	
Novatel_U720	1410	2110		evdo_us	ttyUSB0	option	default (100)		Modify	Delete
Novatel_U727	1410	4100		evdo_us	ttyUSB0	option	default (100)		Modify	Delete
Kyocera_KPC680	0C88	180A		evdo_us	ttyUSB0	option	default (100)		Modify	Delete
Sierra_Compass_597	1199	0023		evdo_us	ttyUSB0	sierra	default (100)		Modify	Delete
Pantech_UM175	106C	3714		evdo_us	ttyUSB1	option	default (100)		Modify	Delete
Sierra_USBCConn_881	1199	6856		gsm_us	ttyUSB0	option	default (100)		Modify	Delete
USBCConn_Mercury_C885	1199	6880		gsm_us	ttyUSB3	option	default (100)		Modify	Delete
Globetrotter_Icon322	0AF0	D033		gsm_us	ttyHS3	hso	default (100)		Modify	Delete

New

Modify Cellular Profile (* indicates required field)

* Cellular Profile Name (max. length: 63) Novatel_U720 * USB Vendor ID (hex) 1410

* USB Product ID (hex) 2110 * USB Serial No. (max. length: 63)

* Dialer Group Setting evdo_us Modem TTY port (max. length: 31) ttyUSB0

Cellular Modem Driver option Modeswitch none eject scsi

Modeswitch Parameter Priority (range: 1-255) default

User Name Authentication (max. length: 63) Password Authentication (max. length: 63)

Retype Password Authentication (max. length: 63)

Modify Cancel

Apply

Commands [View Commands](#)

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 80](#)), 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 80 *Configuring Dialer Group*

The screenshot shows the ArubaOS 5.0 configuration interface. The top navigation bar includes tabs for Monitoring, Configuration (selected), Diagnostics, Maintenance, Plan, Events, and Reports. A 'Save Configuration' button and a 'Logout admin' link are also visible. The left sidebar contains a tree view with categories: Wizards (AP Wizard, Controller Wizard, LAN Wizard, License Wizard), Network (Controller, VLANs, Ports, Uplink (selected)), Security (Authentication, Access Control), Wireless (AP Configuration, AP Installation), and Management (General, Administration, Certificates, SNMP, Logging, Clock). The main content area is titled 'Network > Uplink > Dialer Group' and has three sub-tabs: Uplink Manager, Cellular Profile, and Dialer Group (selected). The 'Dialer Group Table' shows two entries: 'evdo_us' with Init String 'ATQ0V1E0' and Dial String 'ATDT#777', and 'gsm_us' with Init String 'AT+CGDCONT=1,"IP","ISP.CINGULAR"' and Dial String 'ATD*99#'. Each entry has 'Modify' and 'Delete' buttons. Below the table is a 'New' button. The 'Modify Dialer Group' section shows the 'evdo_us' group being edited, with input fields for Name, Init String (ATQ0V1E0), and Dial String (ATDT#777), and 'Modify' and 'Cancel' buttons. An 'Apply' button is at the bottom right. A 'Commands' section at the bottom has a 'View Commands' link.

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. Verify that the modem is detected (`show usb` command)

Figure 81 *Display supported USB modems*

```
(host) #show usb
Address  Product                Vendor  ProdID  Serial                Type      Profile      State
-----  -
3        Novatel Wireless CDMA  1410   4100    091087843891000     Cellular  Novatel_U727 Device ready

(host) #
```

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.

Figure 82 *show usb verbose example (partial)*

```
(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
...
(host) #
```

3. Verify the modem is registered with the Uplink Manager.

Figure 83 *show uplink*

```
(host) #show uplink
Id  Uplink Type  Properties  Priority  State  Status
----
1   Wired      vlan 1     200      Connected * Active *
2   Cellular  Novatel_U727 100      Standby  Ready
(host) #
```

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.

Figure 84 *uplink cellular priority*

```
(host) (config) #uplink cellular priority 201
(host) (config) #
```

4. 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.
5. 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 Controller.



Note: 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 85](#))

Figure 85 *show usb command*

```
(host) #show usb
Address  Product  Vendor  ProdID  Serial  Type  Profile  State
-----
3        Novatel Wireless CDMA 1410 4100 091087843891000 Cellular Novatel_U727 Device ready
(host) #
```

If your modem is not recognized (such as “type is unknown”, “no matching profile”, or “device not ready”), use the `show usb verbose` ([Figure 86](#)) command to verify your modem is listed.

Figure 86 *show usb verbose for profile and driver*

```
(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
...
(host) #
```

3. Create a cellular profile and import the identifiers. The Dialer, Tty, and Driver fields are the new profile defaults.

Figure 87 *cellular profile new_card command*

```
(host) (config) #cellular profile new_card
(host) (config-cellular new_card) # import 10
(host) (config-cellular new_card) # show cellular profile
```

Cellular Profile Table

Name	Vend	Prod	Serial	Dialer	Tty	Driver	Priority	Modeswitch
new_card	1410	5010	091087843890000	evdo_us	ttyUSB0	option	default	

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

4. 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. The following example has four serial TTY ports (option driver) and one flash device (usb-storage driver).

Figure 88 *Driver options*

```
(host) #show usb verbose
...
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
I: If#= 0 Alt= 0 #EPs= 3 Cls=ff(vend.) Sub=ff Prot=ff Driver=option
I: If#= 1 Alt= 0 #EPs= 2 Cls=ff(vend.) Sub=ff Prot=ff Driver=option
I: If#= 2 Alt= 0 #EPs= 2 Cls=ff(vend.) Sub=ff Prot=ff Driver=option
I: If#= 3 Alt= 0 #EPs= 2 Cls=ff(vend.) Sub=ff Prot=ff Driver=option
I: If#= 4 Alt= 0 #EPs= 2 Cls=08(stor.) Sub=06 Prot=50 Driver=usb-storage
...

```

If you get entries similar to the example below:

Figure 89 *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 [“Mode-Switching” on page 428](#)).

Configuring the TTY Port

1. View the exposed TTY ports by executing the `show usb ports 13` command.

Figure 90 *show usb ports 13 command*

```
(host) (config-cellular new_card)# show usb ports 13
ttyUSB0
ttyUSB1
ttyUSB2
ttyUSB3
(host) (config-cellular new_card)#

```

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 [“Mode-Switching” on page 428](#) for instruction.

2. Send a test AT command to determine the correct modem port.

Figure 91 *show usb test command*

```
(host) (support)#show usb test 16 ttyUSB0
AT
OK
TTY port responded to modem AT commands
(host) (support)#

```

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 92 *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 93 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 94 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  Device ready
(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.



Note: 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 95 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
OK
TTY port responded to modem AT commands
(host) #
```

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 96](#))

Figure 96 *show dialer group example*

```
(host)# show dialer group
Dialer Group Table
-----
Name          Init String          Dial String
-----
evdo_us       ATQOV1E0             ATDT#777
gsm_us        AT+CGDCONT=1,"IP","ISP.CINGULAR" ATD*99#
(host)#
```

The ATD, in the Dial String column in [Figure 96](#), 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:

- <http://www.evdoforums.com/>
- <http://ubuntuforums.org>
- <http://www.linux.com/forums>
- <http://kenkinder.com/>

NAS (Network-Attached Storage)

The Dell PowerConnect W-600 Controller Series controller allows you to connect a pre-formatted NAS device that can be made available to all connected clients. The W-600 Series Controller supports NAS devices with partitions in filesystem formats:

- ext2
- ext3
- FAT16
- FAT32

The W-600 Series Controller supports a maximum of four devices. To ensure higher reliability, only connect one USB powered device. The other three devices should use an external power source.

NAS Device Setup

Setting up a NAS device involves the following step:

- Connecting the physical device to the USB port in the controller
- Mounting the device on the controller
- Creating a share—To use the mounted NAS device, you must create a share on the NAS device.
- Associating the share with a filesystem path

Power on the NAS device after you connect the NAS device to the W-600 Series Controller's USB port. Verify that the usb disk is detected (show usb command).

```
(host) #show usb
USB Device Table
-----
Address  Product          Vendor  ProdID  Serial          Type    Profile  State
-----  -
5        OneTouch         0d49   7350    2HAS49ZZ        Storage
3                                     0424   2502                                Hub
4        HP LaserJet P3005 03f0   7317    CNH1D00105      Printer
```

Configuring in the CLI

1. Login as admin and switch to config mode.
2. Enter the command below to enable NAS service:

```
(host) (config)# service network-storage
```

3. Enter the **show usb-storage** command to view a list of mounted and unmounted devices:

```
(host) (config) #show usb-storage
USB Disk Table
-----
Device Name                Device Alias  Num of Partitions  Size      Mounted partitions
-----
Maxtor-Basics_Desktop-2HBADMJ4  Maxtor1TB    1                  1000 GB   No
WD-2500BEV_External-WD-WXE508ET3777  WD250GB     1                  250 GB    No
```

4. Enter the **show usb-storage partitions** command to view disk partitions:

```
(host) (config) #show usb-storage partitions
USB Disk Partition Table
-----
Partition Name                Partition Alias  Filesystem  Size  Used  Mount Name
-----
Maxtor-Basics_Desktop-2HBADMJ4_p1  MxDocs         EXT3/EXT2  1000  204.2M  Maxtor-Basics_Desktop-2HBADMJ4_p1
WD-2500BEV_External-WD-WXE508ET3777_p1  WdImages       EXT3/EXT2  250   223.1M  WD-2500BEV_External-WD-WXE508ET3777_p1
```

5. Enter the command below to create a share:

```
(host) (config)# network-storage share <sharename>
```

6. Associating the share to a filesystem path—To access the share, you must create a filesystem path to the share. enter:

```
(host) (config-network-storage share)# share usb: disk <disk name> <filesystem path>
mode
```

Where,

disk name is the name of the disk. You can also specify the disk alias instead of the disk name.

filesystem path is the path to access the share. This path contains the partition name and the shared folder name.

mode is the permission settings. You can either specify read-only or read-write modes.

Example: share usb: disk WD250GB WdImages/desktop mode Read-Write

7. Display the status of a connected NAS device, enter the command:

```
(host) (config)# show network-storage status
```

Users can now access the connected storage device from the filesystem path.

For example: `\\<controller-ip>\<sharename>\<directory>`

Managing NAS Devices

The following commands are available for managing a NAS devices after they are mounted and configured in the controller. For more details on these command, see the *Dell PowerConnect Command Line Reference Guide*.

- Creating an alias for a disk
`usb-storage disk WD-2500BEV_External-WD-WXE508ET3777 alias WD250GB`
- View list of shares in a disk
`show network-storage shares`
Displays the disk name, partition name, folder and share name, share path, permission settings and status.
- View list of files opened by clients
`show network-storage files opened`
Displays the client machine IP address, path to opened file in controller, permission settings and time-stamp details.
- View list of connected users
`show network-storage users`
Displays the list of users by IP address, connected share name and connection time.
- View list of directories in a disk
`show dir usb: disk <disk-name> <filesystem-path>`
Displays the list of directories in the specified disk and the filesystem path.
- View mounted and unmounted storage device status
`show usb-storage`
Displays device name, device alias (if any), number of partitions in the device, size and mounted partition status of all disks connected to the controller.
- View mounted storage device status (see
`show usb-storage mounted`
- View unmounted storage device status
`show usb-storage unmounted`
Displays if the partitions in the connected disks are unmounted.
- View details of both mounted and unmounted disk partitions
`show usb-storage partitions`
- View details of unmounted disk partitions
`show usb-storage unmounted partitions`
- View details of mounted disk partitions
`show usb-storage mounted partitions`

Mounting and Unmounting Devices

Users who don't have access to the CLI/WebUI can unmount/mount all the disks using the media eject button. This multi-function button means that pressing and holding the button for shorter or longer periods of time will result in entirely different functions. [Table 88](#) list the functions and related status LED for the multi-function eject button.

Table 88 Multi-function Media Eject Button


Initial State	LED State	Action	Status LED	Function	LED Action Completed
NAS Media Operational	Green-solid	Press and hold media eject button for 1 to 5 seconds only	Amber-flashing	Un-mount all NAS media	Amber-solid
NAS Media Unmounted	Amber-solid	Press and hold media eject button for 1 to 5 seconds only	Amber-flashing	Mount all attached NAS devices, and return to fully functional operation	Green-solid
Operational	Green-solid	Press and hold media eject button for more than 5 seconds only	Red-flashing	Controller goes into Standby	Red-solid
Operating with NAS Media un-mounted	Amber-solid	Press and hold media eject button for more than 5 seconds only	Red-flashing	Controller goes into Standby	Red-solid
Standby	Red-solid	Press media eject button	Amber-flashing	Controller wake-up	Green-solid

In the WebUI

You can set up and configure a NAS device using the W-600 Series Controller's WebUI.


The NAS management options are available in the Configuration tab of the WebUI. Go to the **Configuration** tab and click **Disks** under **Management**. This will display the list of connected NAS devices and the clients using the device.

If the NAS service is not enabled, a blank page with a message and hyperlink to enable the NAS service is displayed.



1. Select the [Click here to enable](#) hyperlink to enable the NAS service. After the NAS service is enabled, a list of mounted devices is displayed.
2. Select the disk name to view more details about the mounted device. The following details about the device are displayed on the page:
 - Partition name
 - Filesystem type
 - Total size of the disk
 - Total used size
 - Mount name of the disk


- You can rename or unmount a disk by right-clicking on the disk.



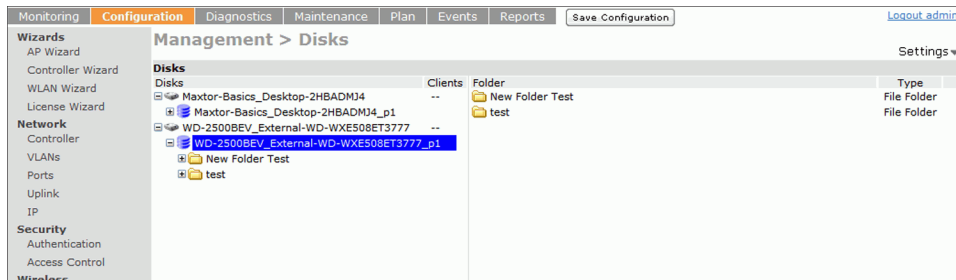
- Rename**—To rename a disk, right click on the disk name and select the **Rename** option. In the pop-up window, enter a new name for the disk and click the **Ok** button.



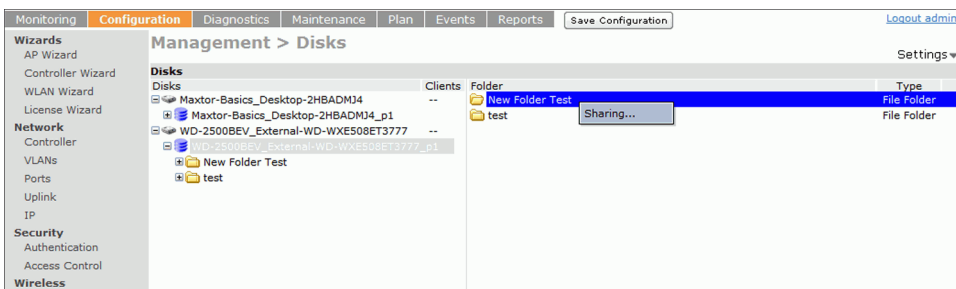
- Un-mount**—To un-mount a disk from the controller, right click on the disk name and select **Unmount** option. In the pop-up window, click the **Unmount** button confirm.




- To view the list of directories in a mounted disk, expand and click on the partition name.



- To share a folder, right click on the folder name and click the **Sharing** button.



6. Sharing folder—To enable share, click Share this folder check box and enter a name for the share. You can also set the access rights for the folder.



Print Server

The W-600 Series Controller allows you to connect a printer so that it is available to all connected clients.

Printer Setup

Connect the printer to the controller's USB port and power on the printer. Then you can configure the printer using either the CLI or the WebUI.

In the CLI

1. Login to the W-600 Series Controller as an admin and switch to config mode.
2. Enable the printer service by entering the command:

```
(host) (config)# service print-server
```
3. To view a list of printers mounted on the controller, type:

```
(host)# show network-printer status
```
4. You can create a printer alias name so that it is identified easily in the network. To create an alias, switch to *config mode* and enter the command:

```
(host) # usb-printer <printer-name> alias <new-printer-name>
```
5. Defining client association
 - a. Maximum clients—You can define the maximum number of clients that can use the printer. Enter the command:

```
(host) (config)# network-printer max-clients <2-20>
```

Currently, the Dell PowerConnect W-600 Controller Series supports a maximum of 20 concurrent clients.
Maximum number of clients per host—To define the maximum number of concurrent clients for a single host, enter the command:

```
(host) (config)# network-printer max-clients-per-host <1-20>
```

Currently, the W-600 Series Controller supports a maximum of 20 concurrent clients.
6. Defining printer job storage—To view the maximum number of jobs that can be saved in the memory, type:

```
(host) (config)# network-printer max-jobs <1-50>
```

Currently, the W-600 Series Controller will support a storage of 50 jobs.

You can now access the printer from their clients.

For example: `\\<controller-ip>\<printername>`

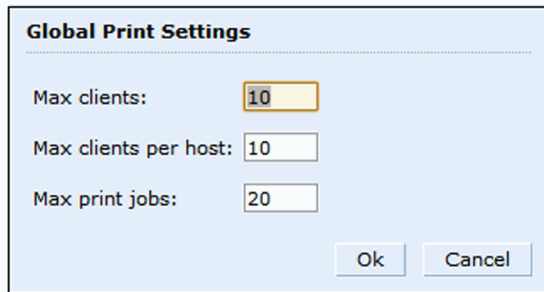
In the WebUI

You can set up and attach a printer using the W-600 Series Controller's WebUI.

The printer management options are available in the Configuration tab of the WebUI. Go to the **Configuration** tab and click **Printer** under **Management**. This will display the list of connected printer and the clients using the printer.

If the printer service is not enabled, a blank page with a message and hyperlink to enable the printer service is displayed.

1. Select the Click here hyperlink to enable the printer service. Once the printer service is enabled, the Printers page displays a list of printers.
2. You can configure global settings for the printer by click the **Settings** hyperlink on the top-right corner of the Printers page.



The image shows a 'Global Print Settings' dialog box with a light blue background. It contains three input fields: 'Max clients' with a value of 10, 'Max clients per host' with a value of 10, and 'Max print jobs' with a value of 20. At the bottom right, there are 'Ok' and 'Cancel' buttons.

The Global Print Settings pop-up window allows you to configure the following details:

- Max clients—Maximum number of the clients that be simultaneously connected to the printer. Maximum allowed is 20.
 - Max clients per host—Maximum number of client per host that can be simultaneously connected to the printer. Maximum allowed is 20.
 - Max print jobs—Maximum number of print jobs that can be stored in the printer memory. Maximum allowed is 50.
3. You can rename an attached printer so that it is easily identified in the network. To rename a printer, right click on the printer name and select the **Rename** option. In the pop-up window, enter the new name and click the **Ok** button.

Other commands for managing printers

The following commands are available for managing a printer after they are configured in the controller.

- View printer configuration

```
show network-printer config
```

Displays configuration parameter and its assigned value.
- View list of jobs in printer memory

```
show network-printer job <printer-name>
```
- Delete print jobs

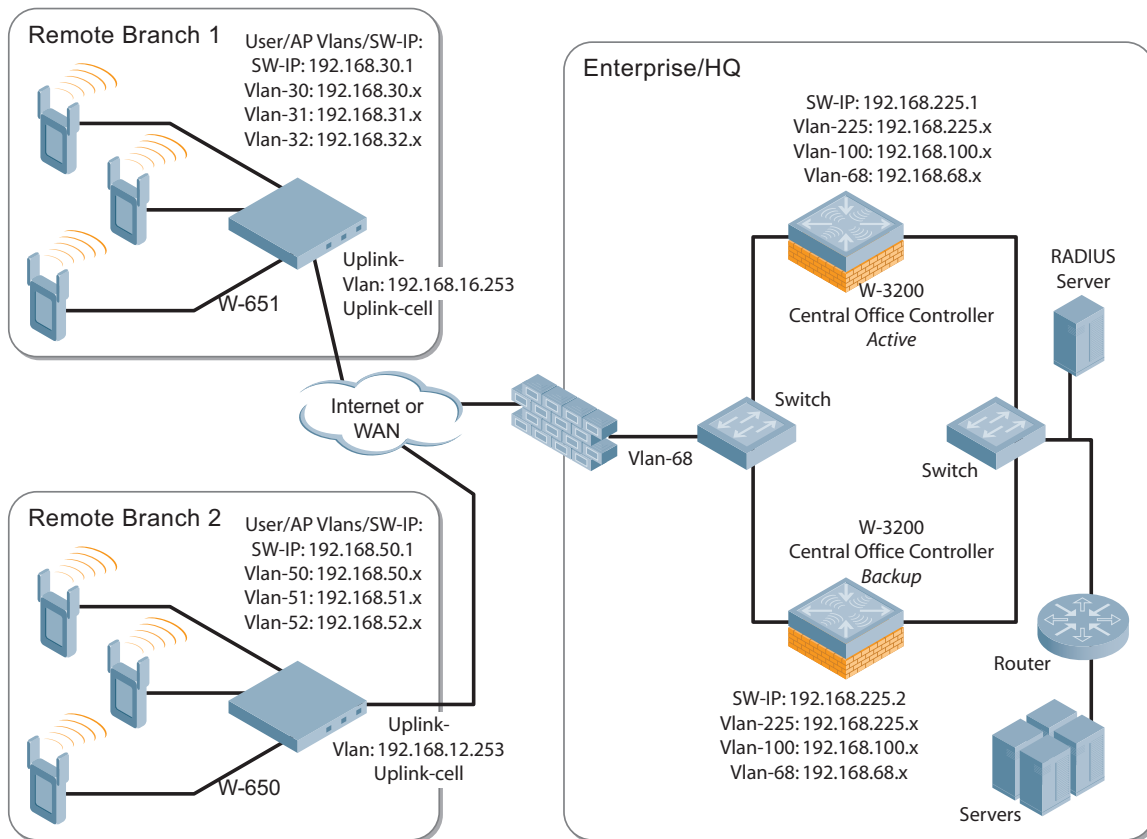
```
network-printer delete <printer-name> job <job-id>
```
- View printer status. The command below displays the printer name, alias, status and status comment.

```
show network-printer status
```

Sample Topology and Configuration

Figure 97 uses both the W-650 and W-651 to illustrate this example topology. Where the W-650 is used, a W-620 could be used just as effectively.

Figure 97 W-650 Sample Topology



Remote Branch 1—W-651 Controller

```

masterip 192.168.68.217 ipsec ***** uplink
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—W-650 Controller

```

masterip 192.168.68.217 ipsec ***** uplink
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 db947e8d1b383813a4070ab0799fa6246b80fc5cfc3268f
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
!
```


OSPFv2 (Open Shortest Path First) is a dynamic Interior Gateway routing Protocol (IGP) based on IETF RFC 2328. The premise of OSPF is that the shortest or fastest routing path is used. Dell's implementation of OSPFv2 allows Dell controllers to be deployed effectively in a Layer 3 topology. Dell controllers can act as default gateway for all clients and forward user packets to the upstream router. The information in this chapter is in the following sections:

- “Important Points to Remember” on page 451
- “WLAN Scenario” on page 451
- “Branch Office Scenario” on page 453
- “Configuring OSPF” on page 455
- “Deployment Best Practices” on page 456
- “Sample Topology and Configuration” on page 457

Important Points to Remember

- OSPF is disabled by default
- Dell controllers support only one OSPF instance
- Maximum OSPF routes is 1K
- Convergence takes between 5 and 15 seconds
- Only stub and totally stub areas are supported
- Only one area can be configured
- An Dell controller can *not* act as ABR (Area border router) or ASBR (Autonomous system border router)
- OSPF packets use generic routing encapsulation (GRE) over Internet Protocol Security (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 an Dell controller is 1100. When running OSPF over a GRE tunnel between an Dell controller and another vendor's router, the MTU values must be the same on both sides of the GRE tunnel.

OSPF is a robust routing protocol addressing various link types and deployment scenarios, the Dell implementation applies to two main use cases; WLAN Scenario and Branch Office Scenario.

WLAN Scenario

In the WLAN scenario, the Dell controller acts as a default gateway for all the clients and talks to one or two (for redundancy) upstream routers. The controller advertises all the user subnet addresses as stub addresses via LSAs to the routers. The controller and upstream routers are part of a totally stub area (TSA). The upstream routers advertise only the default route to the controller.




Note: Totally stub areas *see* only a default route and routes local to the areas themselves.

WLAN Topology

The controller (Figure 98) 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

Figure 98 WLAN OSPF Topology



Based on the cost of the uplink interface, 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

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*
C     4.1.1.0 is directly connected, VLAN4
C     5.1.1.0 is directly connected, VLAN5
C    10.1.1.0 is directly connected, VLAN10
C    12.0.1.0 is directly connected, VLAN12
```

Below is the routing table for Router 1:

```
(router1) #show ip route

Codes: C - connected, O - OSPF, R - RIP, S - static
       M - mgmt, U - route usable, * - candidate default

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

Codes: C - connected, O - OSPF, R - RIP, S - static
       M - mgmt, U - route usable, * - candidate default

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
```


Branch Office Scenario

The branch office scenario has a number of remote branch offices with controllers talking to a central office via an Dell 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 99](#)). 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 99 Branch Office OSPF Topology



All the branch office controllers, the Central office controller, and the upstream routers are part of a stub area. Since 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 interoperate with other vendors seamlessly.

In [Figure 99](#), 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

Gateway of last resort is 20.1.1.2 to network 0.0.0.0

O*    30.0.0.0/0 [1/0] via 20.1.1.2*
C     14.1.1.0 is directly connected, VLAN14
C     15.1.1.0 is directly connected, VLAN15
C     20.1.1.0 is directly connected, Tunnel 1
```

The routing table of the Central office controller is below:

```
(host) #show ip route

Codes: C - connected, O - OSPF, R - RIP, S - static
       M - mgmt, U - route usable, * - candidate default

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

Codes: C - connected, O - OSPF, R - RIP, S - static
       M - mgmt, U - route usable, * - candidate default

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 Router 2 is below:

```
(router2) #show ip route

Codes: C - connected, O - OSPF, R - RIP, S - static
       M - mgmt, U - route usable, * - candidate default

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

Configure general OSPF settings from the OSPF tab on the **Configuration > IP** page (see [Figure 100](#)). 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 100 General OSPF Configuration

The screenshot shows the 'Network > IP > OSPF' configuration page. The left sidebar contains navigation menus for Wizards, Network, Security, Wireless, and Management. The main content area has tabs for IP Interfaces, IP Routes, GRE Tunnels, DHCP Server, OSPF, and IP IGMP. The OSPF configuration includes:

- Enable OSPF:** A checkbox that is currently unchecked.
- Router ID:** A text field containing '10.3.29.99'.
- Redistribute VLAN ID:** An empty text field.
- Area:** A table with columns for Network, Stub, and No-summary.

	Network	Stub	No-summary
Add	192.168.21	Yes	Yes
- Excluded subnet:** A table with columns for Network and Subnet.

	Network	Subnet
Add	192.165.1.0	255.255.255.0
- Commands:** A section for entering configuration commands.

Configure the OSPF interface settings in the Configuration screen ([Figure 101](#)). If OSPF is enable, the parameters contain the correct default values. The OSPF values are editable only when OSPF is enabled on the interface.

Figure 101 Edit OSPF VLAN Settings

The screenshot shows the 'Network > IP > IP Interface > Edit VLAN (1)' configuration page. The left sidebar is similar to Figure 100. The main content area is divided into several sections:

- Details:** Fields for VLAN ID (1), IP address options (DHCP, Client ID, PPPoE), Service name, Username, Password, Confirm Password, and IP address (10.3.29.99) with Net Mask (255.255.255.0).
- DHCP Helper Addresses:** A section for adding DHCP helper addresses.
- IGMP:** Options for Enable IGMP, Snooping, and Proxy (Interface: Gigabitethernet 1/8).
- NAT:** Option for Enable source NAT for this VLAN.
- Inter-VLAN Routing:** Option for Enable Inter-VLAN Routing (checked).
- MLD:** Options for Enable MLD and Snooping.
- OSPF:** A highlighted section for OSPF configuration:
 - Enable OSPF:
 - Area Network (eg. 192.168.1.1): [Empty field]
 - Authentication: Message-digest
 - Message-digest Key: Key [1-255] 1, Password [Empty field]
 - Cost [1-65535]: 1
 - Dead Interval [1-65535]: 40
 - Hello Interval [1-65535]: 10
 - Priority [0-255]: 1
 - Retransmit Interval [1-65535]: 5
 - Transmit Delay [1-65535]: 1

OSPF monitoring is available from an IP Routing sub-section (see [Figure 102](#)). Both Static and OSPF routes are available in table format.

Figure 102 *Monitoring OSPF*

The screenshot shows the ArubaOS 5.0 monitoring interface. The top navigation bar includes tabs for Monitoring, Configuration, Diagnostics, Maintenance, Plan, Events, and Reports. The left sidebar lists various network and controller components, with 'IP Routing' selected. The main content area is titled 'Controller > IP Routing > Routing' and contains two tabs: 'Routing' and 'OSPF'. Below the tabs is a table titled 'Static/OSPF Routes' with the following data:

Type	Network	Via
static	0.0.0.0	10.3.29.254

OSPF Interfaces and Neighboring information is available from the OSPF tab (see [Figure 102](#)). The Interface information includes transmit (TX) and receive (RX) statistics.

Deployment Best Practices

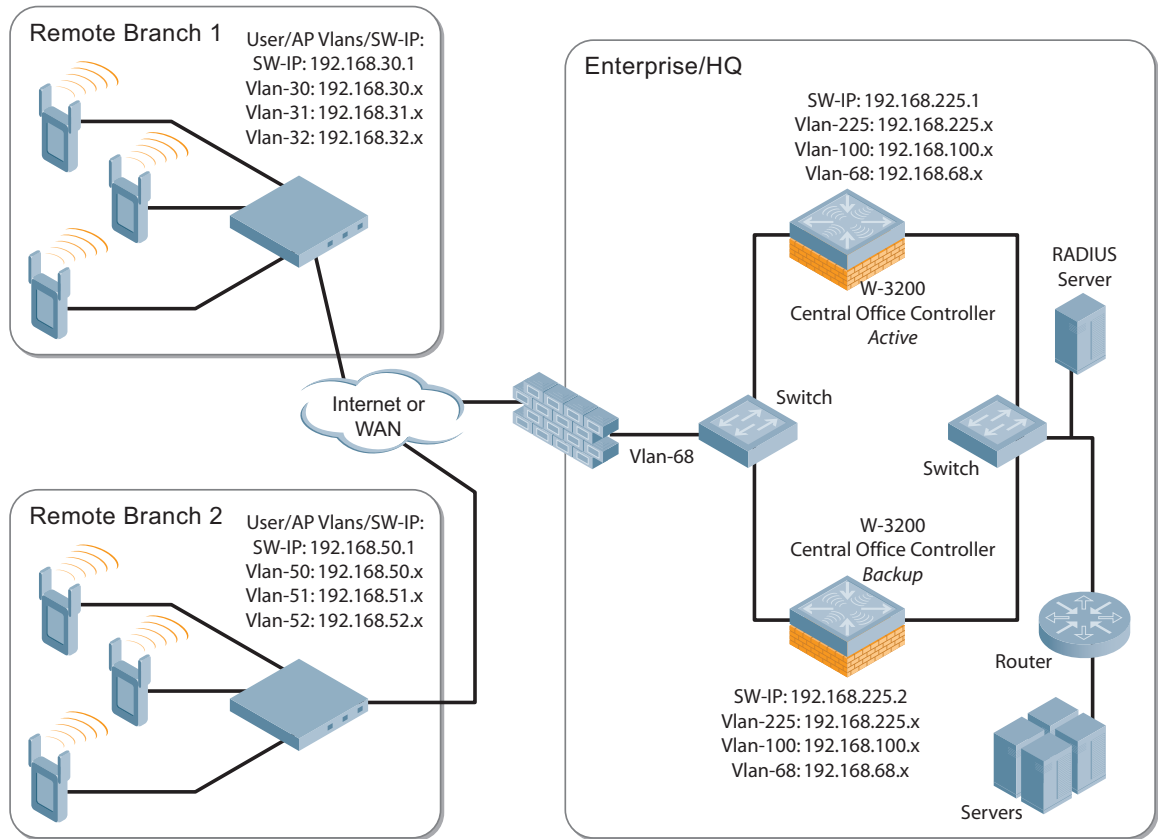
Below are some guidelines regarding deployment and topology for this release of OSPFv2.

- In WLAN scenario, configure the Dell controller and all upstream routers in totally stub area; in 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 GRE tunnel must be in Layer 2 mode to support Multicast over IPsec tunnel. The Corporate VLAN and the Remote Office VLAN must be the same VLAN ID.
- The default MTU value for a Layer 3 GRE tunnel in an Dell controller is 1100. When running OSPF over a GRE tunnel between an Dell controller and another vendor's router, the MTU values must be the same on both sides of the GRE tunnel.
- The controllers do not support ABR/ASBR; designate the upstream router (or some other router) as a ABR.
- 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.

Sample Topology and Configuration

Figure 103 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 103 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 db947e8d1b383813a4070ab0799fa6246b80fc5cfc3268f
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
!
```

This chapter describes how to configure various intrusion detection system (IDS) capabilities of the Dell user-centric network. The Dell network offers a variety of IDS/intrusion prevention system (IPS) features that you can configure and deploy as required. Like most other security-related features of the Dell network, the IDS configuration is done completely on the master controller in the network.



Note: To use many of the IDS features described in this chapter, you must install a Wireless Intrusion Protection (WIP) license on all controllers in your network. If you install a WIP license on a master controller only, an AP or AM terminated on a local controller will not provide IDS features.

This chapter describes the following topics:

- “IDS Features” on page 463
- “IDS Configuration” on page 466
- “WLAN Management System” on page 480
- “Client Blacklisting” on page 484

IDS Features

This section describes IDS features provided by the Dell system.

Unauthorized Device Detection

Unauthorized device detection includes the ability to detect and disable rogue APs and other devices that can potentially disrupt network operations.

Rogue/Interfering AP Detection

The most important IDS functionality offered in the Dell system is the ability to classify an AP as either a *rogue* AP or an *interfering* AP. 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.

You can enable a policy to automatically disable APs that are classified as rogue APs by the Dell system. When a rogue AP is disabled, no wireless stations are allowed to associate to that AP. Refer to “Configuring the Unauthorized Device Profile via the WebUI” on page 479 and “[Configuring the Unauthorized Device Profile via the CLI](#)” on page 479 for details on how to configure rogue AP detection, classification, and containment.



Note: Rogue AP detection and containment are available in the base operating system.

You can manually reclassify an interfering AP. Refer to “[Classifying APs](#)” on page 482 for details on how to change the classification of an AP.

Adhoc Network Detection and Containment

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. The Dell system can perform both ad-hoc network detection and also disable ad-hoc networks when they are found.

Wireless Bridge Detection

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.

Misconfigured AP Detection

If desired, a list of parameters can be configured that defines the characteristics of a valid AP. This is primarily used when non-Dell APs are being used in the network since the Dell controller cannot configure the third-party APs. These parameters can include preamble type, WEP configuration, OUI of valid MAC addresses, valid channels, DCF/PCF configuration, and ESSID. The system can also be configured to detect an AP using a weak WEP key. If a valid AP is detected as misconfigured, the system will deny access to the misconfigured AP if protection is enabled. In cases where someone gains configuration access to a third-party AP and changes the configuration, this policy is useful in blocking access to that AP until the configuration can be fixed.

Weak WEP Detection

The primary means of cracking WEP keys is by capturing 802.11 frames over an extended period of time and searching for patterns of WEP initialization vectors (IVs) that are known to be weak. The Dell system will monitor for devices using weak WEP implementations and generate reports for the administrator of which devices require upgrades.

Multi Tenancy Protection

The Dell system provides the ability to configure SSID lists, and disable unrecognized APs using these reserved resources. This feature can be used in a multi-tenant building where different enterprises must share the RF environment. This feature can also be used to defend against “honeypot” APs. A “honeypot” AP is an attacker’s AP that is set up in close proximity to an enterprise, advertising the ESSID of the enterprise. The goal of such an attack is to lure valid clients to associate to the honeypot AP. From that point, a man in the middle (MITM) attack can be mounted, or an attempt can be made to learn the client’s authentication credentials. Most client devices have no way of distinguishing between a valid AP and an invalid one – the devices only look for a particular ESSID and will associate to the nearest AP advertising that ESSID.

MAC OUI Checking

The Dell system provides the ability to match MAC addresses seen in the air with known manufacturers. The first three bytes of a MAC address are known as the MAC OUI (Organizationally Unique Identifier) and are assigned by the IEEE. Often, clients using a spoofed MAC address will not use a valid OUI, and instead use a randomly generated MAC address. By enabling MAC OUI checking, administrators will be notified if an unrecognized MAC address is in use.

Denial of Service (DoS) Detection

DoS attacks are designed to prevent or inhibit legitimate clients from accessing the network. This includes blocking network access completely, degrading network service, and increasing processing load on clients and

network equipment. Denial of Service attack detection encompasses both rate analysis and the detection of a specific DoS attack known as Fake AP.

Rate Analysis

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. The Dell controller can be configured with the thresholds that indicate a DoS attack and can detect the same. Refer to [“Denial of Service \(DoS\) Detection” on page 464](#) for more details.

Fake AP

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 different APs in the area, thus concealing the real AP. While the tool is still effective for this purpose, a newer purpose is to flood public hotspots or enterprises with fake AP beacons to confuse legitimate clients and to increase the amount of processing client operating systems must do. Refer to [“Denial of Service Attack Detection” on page 468](#) for more details.

Impersonation Detection

A successful man-in-the-middle attack will insert an attacker into the data path between the client and the AP. In such a position, the attacker can delete, add, or modify data, provided he has access to the encryption keys. Such an attack also enables other attacks that can learn a client’s authentication credentials. Man-in-the-middle attacks often rely on a number of different vulnerabilities.

Station Disconnection

Spoofed deauthenticate frames form the basis for most denial of service attacks, as well as the basis for many other attacks such as man-in-the-middle. In a station disconnection attack, an attacker spoofs the MAC address of either an active client or an active AP. The attacker then sends *deauthenticate* frames to the target device, causing it to lose its active association. In addition to a deauthentication frame, Reassociate, Authenticate, and Disassociate frames can also cause the target device to lose its active association.

EAP Handshake Analysis

EAP (Extensible Authentication Protocol) is a component of 802.1x used for authentication. Some attacks, such as “ASLEAP” (used to attack Cisco LEAP) send spoofed deauthenticate messages to clients in order to force the client to re-authenticate multiple times. These attacks then capture the authentication frames for offline analysis. EAP Handshake Analysis detects a client performing an abnormal number of authentication procedures and generates an alarm when this condition is detected.

Sequence Number Analysis

During an impersonation attack, the attacker will generally spoof the MAC address of a client or AP. If two devices are active on the network with the same MAC address, their 802.11 sequence numbers will not match – since the sequence number is usually generated by the NIC firmware, even a custom driver will not generally be able to modify these numbers. Sequence number analysis will detect possible impersonation attacks by looking for anomalies between sequence numbers seen in frames in the air.

AP Impersonation

AP impersonation attacks can be done for several purposes, including as a Man-In-the-Middle attack, as a rogue AP attempting to bypass detection, and as a possible honeypot attack. In such an attack, the attacker sets up an AP that assumes the BSSID and ESSID of a valid AP.

Signature Detection

Many WLAN intrusion and attack tools generate characteristic signatures that can be detected by the Dell network. The system is pre-configured with several known signatures, and also includes the ability for you to create new signatures. For more details on how to configure and create new signatures refer to “[Signature Detection](#)” on page 466.

IDS Configuration

This section describes how to configure IDS features using the IDS profiles. You apply the top-level IDS profile to an AP group or specific AP.

IDS Profile Hierarchy

The top-level IDS profile, assigned to an AP group or AP name, refers to the following IDS profiles:

Table 89 *IDS Profiles*

Profile	Description
IDS General profile	Configures AP attributes.
IDS Rate Thresholds profile	Defines thresholds assigned to the different frame types for rate anomaly checking.
IDS Signature Matching	Configures signatures for intrusion detection. This profile can include predefined signatures or signatures that you configure.
IDS DoS profile	Configures traffic anomalies for Denial of Service attacks.
IDS Impersonation profile	Configures anomalies for impersonation attacks.
IDS Unauthorized Device profile	Configures detection for unauthorized devices. Also configures rogue AP detection and containment.

ArubaOS includes predefined top-level IDS profiles that provide different levels of sensitivity. The following are predefined IDS profiles:

- ids-disabled
- ids-high-setting
- ids-low-setting (the default setting)
- ids-medium-setting



Note: A predefined IDS profile refers to specific instances of the other IDS profiles. You cannot create new instances of a profile within a predefined IDS profile. You can modify parameters within the other IDS profiles.

Configuring IDS via the WebUI

1. Navigate to the **Configuration > AP Configuration** page. Select either **AP Group** or **AP Specific**.
 - If you selected **AP Group**, click **Edit** for the AP group name for which you want to configure IDS.
 - If you selected **AP Specific**, select the name of the AP for which you want to configure IDS.
2. In the Profiles list, expand the **IDS** menu. Select **IDS profile** to display the IDS profiles that are contained in the top-level profile. You can select a predefined IDS profile or create a new profile.
3. Click **Apply**.

Configuring IDS via the CLI

To configure the IDS profile via the command-line interface, access the CLI in config mode and issue the following commands:

```
ap-group <group>  
  ids-profile <profile>
```

IDS General Profile

Table 90 describes the parameters you can configure in the IDS general profile.

Table 90 *IDS General Profile Configuration Parameters*

Parameter	Description
Stats Update Interval	Time interval, in seconds, for the AP to update the controller with statistics. NOTE: This setting takes effect only if the Dell Mobility Manager is configured. Otherwise, statistics update to the controller is disabled. Default: 60 seconds
AP Inactivity Timeout	Time, in seconds, after which an AP is aged out. Default: 5 seconds
STA Inactivity Timeout	Time, in seconds, after which a STA is aged out. Default: 60 seconds
Min Potential AP Beacon Rate	Minimum beacon rate acceptable from a potential AP, in percentage of the advertised beacon interval. Default: 25%
Min Potential AP Monitor Time	Minimum time, in seconds, a potential AP has to be up before it is classified as a real AP. Default: 2 seconds
Signature Quiet Time	Time to wait, in seconds, after detecting a signature match after which the check can be resumed. Default: 900 seconds
Wireless Containment	Enable/disable containment from the wireless side. Default: enabled
Debug Wireless Containment	Enable/disable debugging of containment from the wireless side. NOTE: Enabling this debug option will cause containment function improperly. Default: disabled
Wired Containment	Enable/disable containment from the wired side. Default: disabled
Mobility Manager RTLS	Enable/disable RTLS communication with the configured mobility-manager. Default: disabled

There are two predefined IDS general profiles, each of which provides different levels of wired and wireless containment. Table 91 describes the settings for each of the predefined profiles:

Table 91 *Predefined IDS General Profiles*

Profile	Wireless Containment	Debug Wireless Containment	Wired Containment
ids-general-disabled	disabled	disabled	disabled
ids-general-high-setting	enabled	disabled	enabled

Configuring the General Profile via the WebUI

1. Navigate to the Configuration > AP Configuration page. Select either AP Group or AP Specific.
 - If you selected AP Group, click Edit for the AP group name for which you want to configure IDS.
 - If you selected AP Specific, select the name of the AP for which you want to configure IDS.
2. Expand the IDS menu. Select IDS profile to display the IDS profiles that are contained in the top-level profile.
3. Select IDS General profile.
4. Select a predefined IDS general profile from the drop-down menu, or modify parameters and click Save As to create a new IDS general profile.



Note: If you selected a predefined IDS profile, you cannot select or create a different IDS general profile instance. You can modify parameters within the IDS general profile instance.

5. Click Apply.

Configuring the General Profile via the CLI

To configure this profile via the command-line interface, access the CLI in config mode and issue the following commands:

```
ids general-profile <profile>
  ap-inactivity-timeout <seconds>
  clone <profile>
  min-pot-ap-beacon-rate <percent>
  min-pot-ap-monitor-time <seconds>
  mobility-manager-rtls
  signature-quiet-time <seconds>
  sta-inactivity-timeout <seconds>
  stats-update-interval <seconds>
  wired-containment
  wireless-containment
  wireless-containment-debug
```

Denial of Service Attack Detection

Table 92 describes the parameters you can configure in the IDS DoS profile.

Table 92 IDS Denial of Service Profile Configuration Parameters

Parameter	Description
Detect Disconnect Station Attack	Enables or disables detection of station disconnection attacks. Default: disabled
Disconnect STA Detection Quiet Time	After a station disconnection attack is detected, the time (in seconds) that must elapse before another identical alarm can be generated. Default: 900 seconds
Detect AP Flood Attack	Enables or disables the detection of flooding with fake AP beacons to confuse legitimate users and to increase the amount of processing need on client operating systems. Default: disabled
AP Flood Threshold	Number of Fake AP beacons that must be received within the Flood Increase Time to trigger an alarm. Default: 50

Table 92 IDS Denial of Service Profile Configuration Parameters (Continued)

Parameter	Description
AP Flood Increase Time	Time, in seconds, during which a configured number of Fake AP beacons must be received to trigger an alarm. Default: 3 seconds
AP Flood Detection Quiet Time	After an alarm has been triggered by a Fake AP flood, the time (in seconds) that must elapse before an identical alarm may be triggered. Default: 900 seconds
Detect EAP Rate Anomaly	Enables or disables Extensible Authentication Protocol (EAP) handshake analysis to detect an abnormal number of authentication procedures on a channel and generates an alarm when this condition is detected. Default: disabled
EAP Rate Threshold	Number of EAP handshakes that must be received within the EAP Rate Time Interval to trigger an alarm. Default: 60
EAP Rate Time Interval	Time, in seconds, during which the configured number of EAP handshakes must be received to trigger an alarm. Default: 3 seconds
EAP Rate Quiet Time	After an alarm has been triggered, the time (in seconds) that must elapse before another identical alarm may be triggered. Default: 900 seconds
Detect Rate Anomalies	Enables or disables detection of rate anomalies. Default: disabled
Detect 802.11n 40MHz Intolerance Setting	Enables or disables detection of 802.11n 40 MHz intolerance setting, which controls whether stations and APs advertising 40 MHz intolerance will be reported. Default: enabled
Client 40MHz Intolerance Detection Quiet Time	Controls the quiet time (when to stop reporting intolerant STAs if they have not been detected), in seconds, for detection of 802.11n 40 MHz intolerance setting. Default: 900 seconds

There are four predefined DoS profiles, each of which provides different levels of detection and containment. [Table 93](#) describes the settings for each of the predefined profiles:

Table 93 Predefined IDS DoS Profiles

Parameter	ids-dos-disabled	ids-dos-low-setting	ids-dos-medium-setting	ids-dos-high-setting
Detect Disconnect Station Attack	disabled	enabled	enabled	disabled
Disconnect STA Detection Quiet Time	900 seconds	900 seconds	900 seconds	900 seconds
Spoofed Death Blacklist	disabled	disabled	disabled	disabled
Detect AP Flood Attack	disabled	disabled	disabled	disabled
AP Flood Threshold	50	50	50	50
AP Flood Increase Time	3 seconds	3 seconds	3 seconds	3 seconds
AP Flood Detection Quiet Time	900 seconds	900 seconds	900 seconds	900 seconds
Detect EAP Rate Anomaly	disabled	disabled	enabled	enabled
EAP Rate Threshold	60	60	30	60
EAP Rate Time Interval	3 seconds	3 seconds	3 seconds	3 seconds

Table 93 *Predefined IDS DoS Profiles*

Parameter	ids-dos-disabled	ids-dos-low-setting	ids-dos-medium-setting	ids-dos-high-setting
EAP Rate Quiet Time	900 seconds	900 seconds	900 seconds	900 seconds
Detect Rate Anomalies	disabled	disabled	disabled	enabled
Detect 802.11n 40 MHz Intolerance Setting	disabled	enabled	enabled	enabled
Client 40 MHz Intolerance Detection Quiet Time	900 seconds	900 seconds	900 seconds	900 seconds
Rate Thresholds for Assoc Frames	default	default	default	default
Rate Thresholds for Disassoc Frames	default	default	default	default
Rate Thresholds for Deauth Frames	default	default	default	default
Rate Thresholds for Probe Request Frames	default	probe-request-response-thresholds	probe-request-response-thresholds	probe-request-response-thresholds
Rate Thresholds for Probe Response Frames	default	probe-request-response-thresholds	probe-request-response-thresholds	probe-request-response-thresholds
Rate Thresholds for Auth Frames	default	default	default	default

Configuring the DoS profile via the WebUI

- Navigate to the **Configuration > AP Configuration** page. Select either **AP Group** or **AP Specific**.
 - If you selected **AP Group**, click **Edit** for the AP group name for which you want to configure IDS.
 - If you selected **AP Specific**, select the name of the AP for which you want to configure IDS.
- Expand the **IDS** menu. Select **IDS profile** to display the IDS profiles that are contained in the top-level profile.
- Select **IDS DoS profile**.
- You can select a predefined profile from the drop-down menu. Or you can modify parameters and click **Save As** to create an IDS DoS profile instance.



Note: If you selected a predefined IDS profile, you cannot select or create a different IDS DoS profile instance. You can modify parameters within the IDS DoS profile instance.

- Click **Apply**.

Configuring the DoS profile via the CLI

To configure this profile via the command-line interface, access the CLI in config mode and issue the following commands:

```
ids dos-profile <profile>
  ap-flood-inc-time <seconds>
  ap-flood-quiet-time <seconds>
  ap-flood-threshold <number>
  assoc-rate-thresholds <number>
  auth-rate-thresholds <number>
  client-ht-40mhz-intol-quiet-time <seconds>
  clone <profile>
  deauth-rate-thresholds <number>
```

```

detect-ap-flood
detect-eap-rate-anomaly
detect-ht-40mhz-intolerance
detect-rate-anomalies
disassoc-rate-thresholds <number>
eap-rate-quiet-time <seconds>
eap-rate-threshold <number>
eap-rate-time-interval <seconds>
probe-request-rate-thresholds <number>
probe-response-rate-thresholds <number>
spoofed-deauth-blacklist

```

IDS Rate Thresholds Profile

IDS rate threshold profile defines thresholds assigned to the different frame types for rate anomaly checking. A profile of this type is attached to each of the following 802.11 frame types in the IDS Denial of Service profile:

- Association frames
- Disassociation frames
- Deauthentication frames
- Probe Request frames
- Probe Response frames
- Authentication frames

A channel threshold applies to an entire channel, while a node threshold applies to a particular client MAC address. Dell provides predefined default IDS rate thresholds profiles for each of these types of frames. Default values depend upon the frame type.

[Table 94](#) describes the parameters you can configure for the IDS rate threshold profile.

Table 94 *IDS Rate Thresholds Profile Configuration Parameters*

Parameter	Description
Channel Increase Time	Time, in seconds, in which the threshold must be exceeded in order to trigger an alarm.
Channel Quiet Time	After an alarm has been triggered, the time that must elapse before another identical alarm may be triggered. This option prevents excessive messages in the log file.
Channel Threshold	Specifies the number of a specific type of frame that must be exceeded within a specific interval in an entire channel to trigger an alarm.
Node Quiet Time	After an alarm has been triggered, the time that must elapse before another identical alarm may be triggered. This option prevents excessive messages in the log file.
Node Threshold	Specifies the number of a specific type of frame that must be exceeded within a specific interval for a particular client MAC address to trigger an alarm.
Node Time Interval	Time, in seconds, in which the threshold must be exceeded in order to trigger an alarm.

Configuring the Rate Threshold Profile via the WebUI

1. In the **Profiles** list, under the IDS DoS profile, select the IDS rate threshold profile you want to configure, or enter the name of a new profile in the entry blank then click **add**.
2. Modify parameters described in [Table 94](#).
3. Click **Apply** to save your changes to the selected profile, or, to create a new IDS rate threshold profile, click **Save As**, enter a new profile name, then click **Apply**.

Configuring the Rate Thresholds Profile via the CLI

To configure this profile via the command-line interface, access the CLI in config mode and issue the following commands:

```
ids rate-thresholds-profile <name>
  channel-inc-time <seconds>
  channel-quiet-time <seconds>
  clone <profile>
  node-quiet-time <seconds>
  node-threshold <number>
  node-time-interval <seconds>
ids dos-profile <profile>
  <frame-type> <thresholds-profile>
```

Impersonation Detection Profile

Table 95 describes the parameters you can configure in the IDS Impersonation profile.

Table 95 IDS Impersonation Profile Configuration Parameters

Parameter	Description
Detect AP Impersonation	Enables or disables detection of 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. Default: enabled
Protect from AP Impersonation	When AP impersonation is detected, both the legitimate and impersonating AP are disabled using a denial of service attack. Default: disabled
Beacon Diff Threshold	Percentage increase in beacon rate that triggers an AP impersonation event. Default: 50%
Beacon Increase Wait Time	Time, in seconds, after the Beacon Diff Threshold is crossed before an AP impersonation event is generated. Default: 3 seconds

Configuring the Impersonation Profile via the WebUI

1. Navigate to the **Configuration > AP Configuration** page. Select either **AP Group** or **AP Specific**.
 - If you selected **AP Group**, click **Edit** for the AP group name for which you want to configure IDS.
 - If you selected **AP Specific**, select the name of the AP for which you want to configure IDS.
2. Expand the **IDS** menu. Select **IDS profile** to display the IDS profiles that are contained in the top-level profile.
3. Select **IDS Impersonation profile**.
4. You can select a predefined profile from the drop-down menu. Or you can modify parameters and click **Save As** to create an IDS impersonation profile instance.



Note: If you selected a predefined IDS profile, you cannot select or create a different IDS impersonation profile instance. You can modify parameters within the IDS impersonation profile instance.

5. Click **Apply**.

Configuring the Impersonation Profile via the CLI

To configure this profile via the command-line interface, access the CLI in config mode and issue the following commands:

```
ids impersonation-profile <profile>
  beacon-diff-threshold <percent>
  beacon-inc-wait-time <seconds>
  clone <profile>
  detect-ap-impersonation
  protect-ap-impersonation
```

Signature Matching Profile

The IDS signature matching profile contains signatures for intrusion detection. This profile can include predefined signatures or signatures that you configure. [Table 96](#) describes the predefined signatures that you can add to the profile.

Table 96 *Predefined Signatures*

Signature	Description
ASLEAP	A tool created for Linux systems that has been used to attack Cisco LEAP authentication protocol.
Null-Probe-Response	An attack with 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.
AirJack	Originally 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 Access Point.
NetStumbler Generic	NetStumbler is a popular wardriving application used to locate 802.11 networks. When used with certain NICs (such as Orinoco), NetStumbler generates a characteristic frame that can be detected.
NetStumbler Version 3.3.0x	Version 3.3.0 of NetStumbler changed the characteristic frame slightly. This signature detects the updated frame.
Deauth-Broadcast	A deauth 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.

Configuring the Signature Matching Profile via the WebUI

1. Navigate to the **Configuration > AP Configuration** page. Select either **AP Group** or **AP Specific**.
 - If you selected **AP Group**, click **Edit** for the AP group name for which you want to configure IDS.
 - If you selected **AP Specific**, select the name of the AP for which you want to configure IDS.
2. Expand the **IDS** menu. Select **IDS profile** to display the IDS profiles that are contained in the top-level profile.
3. Select **IDS Signature Matching profile**.
4. You can select a predefined profile from the drop-down menu. Or you can modify parameters and click **Save As** to create an IDS signature-matching profile instance.



Note: If you selected a predefined IDS profile, you cannot select or create a different IDS signature-matching profile instance. You can modify parameters within the IDS signature-matching profile instance.

5. Click **Apply**.

Configuring the Signature Matching Profile via the CLI

To configure this profile via the command-line interface, access the CLI in config mode and issue the following commands:

```
ids signature-matching-profile <profile>
  clone <profile>
  signature <profile>
```

IDS Signature Profile

Signature rules match an attribute to a value. For example, you can add a rule that matches the BSSID to the value 00:00:00:00:00:0a. [Table 97](#) describes the attributes and values you can configure for a signature rule.

Table 97 Signature Rule Attributes

Attribute	Description
BSSID	BSSID field in the 802.11 frame header.
Destination MAC address	Destination MAC address in 802.11 frame header.
Frame Type	Type of 802.11 frame. For each type of frame further details can be specified to filter and detect only the required frames. It can be one of the following: <ul style="list-style-type: none">• association• auth• beacon• control (all control frames)• data (all data frames)• deauth• disassoc• management (all management frames)• probe-request• probe-response
SSID	For beacon, probe-request, and probe-response frame types, specify the SSID as either a string or hex pattern.
SSID-length	For beacon, probe-request, and probe-response frame types, specify the SSID length. Maximum length is 32 bytes.
Payload	Pattern at a fixed offset in the payload of a 802.11 frame. Specify the pattern to be matched as a string or hex pattern. Maximum length is 32 bytes.
Offset	When a payload pattern is configured, specify the offset in the payload where the pattern is expected to be found in the frame.
Sequence Number	Sequence number of the frame.
Source MAC address	Source MAC address of the 802.11 frame.

Creating a New Signature via the WebUI

1. Navigate to the **Configuration > Advanced Services > All Profiles** page.
2. Expand the **IDS** menu.
3. Scroll the list of profiles to select **IDS Signature Profile**. Enter the name of the new signature profile and click **Add**.
4. Select the new signature profile name to display profile details.
5. Click **New** to add a rule to the profile.
6. After completing configuring the rule to be added, click **Add** to add the rule.
7. Click **Apply**.

Creating a New Signature via the CLI

To configure this profile via the command-line interface, access the CLI in config mode and issue the following commands:

```
ids signature-profile <profile>
  bssid <macaddr>
  clone <profile>
  dst-mac <macaddr>
  frame-type {assoc|auth|beacon|control|data|deauth|disassoc|mgmt|probe-request|
  probe-response} [ssid <ssid>] [ssid-length <bytes>]
  payload <pattern> [offset <number>]
  seq-num <number>
  src-mac <macaddr>
```

Unauthorized Device Detection

Table 98 describes the parameters (and their defaults) you can configure in the IDS unauthorized device detection profile. There is also one default profile and three predefined unauthorized device profiles, each of which provides different levels of detection and containment, as described in “[Default and Predefined IDS Unauthorized Device Profiles](#)” on page 478.

Table 98 *IDS Unauthorized Device Profile Configuration Parameters*

Parameter	Description
Detect Adhoc Networks	Enable or disable detection of adhoc networks. Default: enabled
Protect from Adhoc Networks	Enable or disable protection from adhoc networks. When adhoc networks are detected, they are disabled using a denial of service attack. Default: disabled
Detect Windows Bridge	Enable or disable detection of Windows station bridging. Default: enabled
Detect Wireless Bridge	Enable or disable detection of wireless bridging. Default: enabled
Detect Devices with an Invalid MAC OUI	Enables or disables the checking of the first three bytes of a MAC address, known as the MAC organizationally unique identifier (OUI), 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. Enabling MAC OUI checking causes an alarm to be triggered if an unrecognized MAC address is in use. Default: disabled
MAC OUI detection Quiet Time	The time, in seconds, that must elapse after an invalid MAC OUI alarm has been triggered before another identical alarm may be triggered. Default: 900 seconds
Adhoc Network detection Quiet Time	The time, in seconds, that must elapse after an adhoc network detection alarm has been triggered before another identical alarm may be triggered. Default: 900 seconds
Wireless Bridge detection Quiet Time	The time, in seconds, that must elapse after a wireless bridging alarm has been triggered before another identical alarm may be triggered. Default: 900 seconds
Rogue AP Classification	Enable or disable rogue AP classification. A rogue AP is one that is unauthorized and plugged into the wired side of the network. Any other AP seen in the RF environment that is not part of the valid enterprise network is considered to be “interfering” — it has the potential to cause RF interference but it is not connected to the wired network and thus does not represent a direct threat. Default: enabled

Table 98 IDS Unauthorized Device Profile Configuration Parameters (Continued)

Parameter	Description
Overlay Rogue Classification	Overlay Rogue Classification is classification through valid/rogue APs. A controller uses the wired-mac table of other valid and rogue APs as equivalents of the wired MACs that it sees on our network. When this match is triggered, it makes a note of the AP that helped in this process, and this info will be displayed as the Helper-AP. By default, Overlay Rogue Classification is disabled in ArubaOS 2.x but enabled in later versions of ArubaOS. Default: enabled
Valid Wired Macs	List of MAC addresses of wired devices in the network, typically gateways or servers.
Rogue Containment	By default, rogue APs are only detected but are not automatically disabled. This option automatically shuts down rogue APs. When this option is enabled, clients attempting to associate to a rogue AP will be disconnected from the rogue AP through a denial of service attack. Default: disabled
Allow Well Known MAC	Allows devices with known MAC addresses to classify rogues APs. Depending on your network, configure one or more of the following options for classifying rogue APs: <ul style="list-style-type: none"> • hsrp—Routers configured for HSRP, a Cisco-proprietary redundancy protocol, with the HSRP MAC OUI 00:00:0c. • iana—Routers using the IANA MAC OUI 00:00:5e. • local-mac—Devices with locally administered MAC addresses starting with 02. • vmware—Devices with any of the following VMWare OUIs: 00:0c:29, 00:05:69, or 00:50:56 • vmware1—Devices with VMware OUI 00:0c:29. • vmware2—Devices with VMware OUI 00:05:69. • vmware3—Devices with VMware OUI 00:50:56. If you modify an existing configuration, the new configuration overrides the original configuration. For example, if you configure allow-well-known-mac hsrp and then configure allow-well-known-mac iana , the original configuration is lost. To add more options to the original configuration, include all of the required options, for example: allow-well-known-mac hsrp iana . NOTE: Use caution when configuring this command. If the neighboring network uses similar routers, those APs might be classified as rogues. If containment is enabled, clients attempting to associate to an AP classified as a rogue are disconnected through a denial of service attack. To clear the well known MACs in the system, issue the following CLI commands pm all controllers: <ol style="list-style-type: none"> 1. clear wms wired-mac This clears all of the learned wired MAC information on the controller. 2. reload This reboots the controller.
Suspected Rogue Containment	Suspected rogue APs are treated as interfering APs, thereby the controller attempts to reclassify them as rogue APs. By default, suspected rogue APs are not automatically contained. In combination with the suspected rogue containment confidence level, this option automatically shuts down suspected rogue APs. When this option is enabled, clients attempting to associate to a suspected rogue AP will be disconnected from the suspected rogue AP through a denial of service attack. Default: disabled
Suspected Rogue Containment Confidence Level	When an AP is classified as a suspected rogue AP, it is assigned a 50% confidence level. If multiple APs trigger the same events that classify the AP as a suspected rogue, the confidence level increases by 5% up to 95%. In combination with suspected rogue containment, this option configures the threshold by which containment should occur. Suspected rogue containment occurs only when the configured confidence level is met. Default: 60%

Table 98 *IDS Unauthorized Device Profile Configuration Parameters (Continued)*

Parameter	Description
Protect Valid Stations	Does not allow valid stations to connect to a non-valid AP (see “Classifying APs” on page 482). Default: disabled
Detect Bad WEP	Enables or disables 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. Default: disabled
Detect Misconfigured AP	Enables or disables detection of misconfigured APs. An AP is classified as misconfigured if it does not meet any of the following configurable parameters: <ul style="list-style-type: none"> • Valid channels • Encryption type • Short preamble • List of valid AP MAC OUIs • Valid SSID list Default: disabled
Protect Misconfigured AP	Enables or disables protection of misconfigured APs. Default: disabled
Protect SSID	Enables or disables use of SSID by only valid APs. Default: disabled
Privacy	Enable or disables encryption as valid AP configuration. Default: disabled
Require WPA	When enabled, any valid AP that is not using WPA encryption is flagged as misconfigured. Default: disabled
Valid 802.11a channel for policy enforcement (multi-valued)	List of valid 802.11a channels that third-party APs are allowed to use. Default: N/A
Valid 802.11g channel for policy enforcement (multi-valued)	List of valid 802.11g channels that third-party APs are allowed to use. Default: N/A
Valid MAC OUIs (multi-valued)	List of valid MAC organizationally unique identifiers (OUIs).
Valid and Protected SSIDs (multi-valued)	List of valid and protected SSIDs.
Protect 802.11n High Throughput Devices	Enables or disables protection of high-throughput (802.11n) devices. Default: disabled.
Protect 40MHz 802.11n High Throughput Devices	Enables or disables protection of high-throughput (802.11n) devices operating in 40 MHz mode. Default: disabled
Detect Active 802.11n Greenfield Mode	Enables or disables detection of high-throughput devices advertising greenfield preamble capability. Default: enabled

The default and predefined IDS unauthorized device profiles are shown in the table below. The Default profile is the equivalent of an “ids-unauthorized-device-low-setting” profile.

Table 99 *Default and Predefined IDS Unauthorized Device Profiles*

Parameter	ids-unauthorized-device-disabled	Default	ids-unauthorized-device-medium-setting	ids-unauthorized-device-high-setting
Detect adhoc networks	disabled	enabled	enabled	enabled
Protect from adhoc networks	disabled	disabled	disabled	enabled
Detect windows bridge	disabled	enabled	enabled	enabled
Detect wireless bridge	disabled	enabled	enabled	enabled
Detect devices with invalid MAC OUI	disabled	disabled	disabled	enabled
MAC OUI detection quiet time	900 seconds	900 seconds	900 seconds	900 seconds
Adhoc network detection quiet time	900 seconds	900 seconds	900 seconds	900 seconds
Wireless bridge detection quiet time	900 seconds	900 seconds	900 seconds	900 seconds
Rogue AP classification	disabled	enabled	enabled	enabled
Overlay rogue AP classification	enabled	enabled	enabled	enabled
Valid wired MACs	—	—	—	—
Allow well known MAC	—	—	—	—
Rogue containment	disabled	disabled	disabled	enabled
Suspected rogue containment	disabled	disabled	disabled	disabled
Suspected rogue containment confidence level	60	60	60	60
Protect valid stations	disabled	disabled	disabled	enabled
Detect bad WEP	disabled	disabled	enabled	enabled
Detect misconfigured AP	disabled	enabled	enabled	enabled
Protect misconfigured AP	disabled	disabled	disabled	enabled
Protect SSID	disabled	disabled	disabled	enabled
Privacy	disabled	disabled	disabled	enabled
Require WPA	disabled	disabled	enabled	disabled
Valid 802.11g channel for policy enforcement	—	—	—	—
Valid 802.11a channel for policy enforcement	—	—	—	—
Valid MAC OUIs	—	—	—	—
Valid and protected SSIDs	—	—	—	—
Protect 802.11n High-throughput Devices	disabled	disabled	disabled	enabled

Table 99 Default and Predefined IDS Unauthorized Device Profiles

Parameter	ids-unauthorized-device-disabled	Default	ids-unauthorized-device-medium-setting	ids-unauthorized-device-high-setting
Protect 40 MHz 802.11n High-throughput Devices	disabled	disabled	disabled	enabled
Detect Active 802.11n Greenfield Mode	disabled	enabled	enabled	enabled

Configuring the Unauthorized Device Profile via the WebUI

1. Navigate to the Configuration > AP Configuration page. Select either AP Group or AP Specific.
 - If you selected AP Group, click Edit for the AP group name for which you want to configure IDS.
 - If you selected AP Specific, select the name of the AP for which you want to configure IDS.
2. Expand the IDS menu. Select IDS profile to display the IDS profiles that are contained in the top-level profile.
3. Select IDS Unauthorized Device profile.
4. You can select a predefined profile from the drop-down menu. Or you can modify parameters and click Save As to create an IDS unauthorized device profile instance.



Note: If you selected a predefined IDS profile, you cannot select or create a different IDS unauthorized device profile instance. You can modify parameters within the IDS unauthorized device profile instance.

5. Click Apply.

Configuring the Unauthorized Device Profile via the CLI

To configure this profile via the command-line interface, access the CLI in config mode and issue the following commands:

```
ids unauthorized-device-profile <profile>
  adhoc-quiet-time <seconds>
  allow-well-known-mac [hsrp|iiana|local-mac|vmware|vmware1|vmware2|vmware3]
  cfg-valid-11a-channel <channel>
  cfg-valid-11g-channel <channel>
  classification
  clone <profile>
  detect-adhoc-network
  detect-bad-wep
  detect-ht-greenfield
  detect-invalid-mac-oui
  detect-misconfigured-ap
  detect-windows-bridge
  detect-wireless-bridge
  mac-oui-quiet-time <seconds>
  overlay-classification
  privacy
  protect-adhoc-network
  protect-high-throughput
  protect-ht-40mhz
  protect-misconfigured-ap
  protect-ssid
  protect-valid-sta
```

```

require-wpa
rogue-containment
suspect-rogue-conf-level <level>
suspect-rogue-containment
valid-and-protected-ssid <ssid>
valid-oui <oui>
valid-wired-mac <macaddr>
wireless-bridge-quiet-time <seconds>

```

WLAN Management System

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.

Configuring WMS via the WebUI

1. Navigate to the Configuration > Advanced Services > Wireless page.
2. Configure the parameters, as described in [Table 100](#).

Table 100 WMS Configuration Parameters

Parameter	Description
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
Mark Known Interfering APs as Persistent Known Interfering APs	Enables or disables APs that are marked as known interfering from being aged out. Default: enabled
Learn APs	Enables or disables AP learning. Learning affects the way APs are classified. Default: disabled

3. Click **Apply**.

Configuring WMS via the CLI

Use the following commands to configure WMS via the CLI. The parameters in this command are described in detail in [Table 100](#).

```
wms general
  ap-ageout-interval <minutes> | collect-stats {disable|enable} |
  learn-ap {enable|disable} | persistent-known-interfering {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.



Note: 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) #wms-local system 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 CLI 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 reinit-db
```

Enabling AP Learning

AP learning is typically used where there are non-Dell APs connected on the same wired network as Dell APs. By default, AP learning is not enabled and any non-Dell APs that are connected on the same networks as Dell APs are classified as rogue APs. Enabling AP learning marks the non-Dell APs as valid APs instead of as rogue APs. You can enable or disable AP learning from the WebUI or CLI.



Note: Enabling AP learning is useful when you install the Dell controller in an environment with an existing third-party wireless network, especially if there are a large number of installed APs. Leave AP learning enabled until all APs in the network have been detected and classified as valid. Then disable AP learning and reclassify any unknown APs as interfering.

To enable AP learning via the WebUI:

1. Navigate to the **Configuration > Advanced Services > Wireless** page.
2. Select (or deselect) the **Learn APs** checkbox.
3. Click **Apply**.

To enable AP learning via the command-line interface, access the CLI in config mode and issue the following commands:

```
wms general learn-ap {enable|disable}
```

Classifying APs

If AP learning is enabled, non-Dell APs connected on the same wired network as Dell APs are classified as valid APs. If AP learning is disabled, a non-Dell AP is classified as a rogue AP. You can also manually classify an AP. For example, if you know about an interfering AP, you can manually reclassify it as a *known* interfering AP. You can use the WebUI or CLI to manually classify an AP into one of the following categories:

Table 101 AP Manual Classification Categories

AP Type	Description
Valid AP	An AP that is part of the enterprise providing WLAN service. Dell APs that successfully connect to the controller and load software and configuration should be classified as valid APs. NOTE: Any client that successfully authenticates with a valid AP and passes encrypted traffic is classified as a valid client. (Encrypted traffic includes encrypted 802.11 frames and unencrypted 802.11 frames which are VPN encrypted.)
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 and is not part of your WLAN network.
Known Interfering AP	An interfering AP where the BSSIDs are known. Once classified, a known interfering AP does not change its state.
Unsecure AP (rogue AP)	A rogue AP is an unauthorized AP that is plugged into the wired side of the network. You can configure automatic shutdown of rogue APs in the IDS unauthorized device detection profile.
Suspected Unsecure AP	A suspected rogue AP is plugged into the wired side of the network, but may not be an unauthorized device. Automatic rogue containment does not apply to suspected rogue APs.
DoS AP	An AP for which denial of service is enabled. Any clients connected to this AP are disconnected.

To enable AP learning via the WebUI:

1. Navigate to the **Reports > AP Reports > All Interfering APs** page on the master controller.
2. Select the checkbox for the AP(s) you want to classify.
3. Click the appropriate **Set as** button on the page.
4. Click **Apply**.

To enable AP learning via the command-line interface, access the CLI in config mode and issue the following commands:

Enter the following command in privilege mode:

```
wms ap <bssid> mode {dos|interfering|known-interfering|unsecure|valid}
```

Configuring Misconfigured AP Detection and Protection

An AP is classified as misconfigured if it does not meet any of the following configurable parameters:

- Valid channels
- Encryption type
- Short preamble

- List of valid AP MAC OUIs
- Valid SSID list (exceptions are described in “Valid Enterprise SSIDs” on page 483)

This classification is primarily for enforcing security policies on non-Dell APs, although the classification and protection mechanism also applies to all valid Dell APs.

Valid Enterprise SSIDs

SSIDs added to the Valid Enterprise SSID list are known as “Valid SSIDs” or “Reserved SSIDs.” The list is empty by default and does not contain any SSIDs configured on the controller. You can add SSIDs to the list using the WebUI or CLI.

To add or remove SSIDs from the Valid Enterprise SSID list via the WebUI:

1. Navigate to the **Configuration > Advanced Services > All Profiles**
2. In the profiles list on the left window pane, expand the **IDS** menu
3. Expand the **IDS Unauthorized Device Profile** menu.
4. In the profiles list, select the name of the profile to which you want to add a valid SSID.
5. Scroll down to the **Valid and Protected SSIDs** section.
 - To add an SSID, Enter the name of the SSID, then click **Add**.
 - To remove an SSID, select an SSID from the list, then click **Delete**.
6. Click **Apply** to save your changes.

To add or remove SSIDs from the list via the command-line interface, access the CLI in config mode and issue the following commands:

```
ids unauthorized-device-profile valid-and-protected-ssid [no] <ssid>
```

As part of its function, Multi-Tenancy protection prevents an interfering AP from advertising an SSID that is added to the Valid Enterprise SSID list. This feature protects against honeypot attacks. Misconfigured AP protection also uses the Valid Enterprise SSID list to classify an AP as misconfigured.

Whether a client can connect to an SSID depends on if a Multi-Tenancy protection or Misconfigured AP protection are enabled or disabled, whether the AP is valid or interfering, and whether the SSID is in the Valid Enterprise SSID list. [Table 102](#) describes client connections to valid and non-valid SSIDs when Multi-Tenancy protection and Misconfigured AP protection are enabled or disabled.

Table 102 *Valid SSIDs with Multi-Tenancy and Misconfigured AP Protection*

Multi-Tenancy Protection	Misconfigured AP Protection	Client Connections
Enabled	Disabled	<p>If there are entries in the valid SSID list:</p> <ul style="list-style-type: none"> • Clients can connect to valid SSIDs on valid APs. • Clients cannot connect to valid SSIDs on interfering APs (including known interfering APs). • Clients can connect to SSIDs not in the valid SSID list on valid APs. • Clients can connect to SSIDs not in the valid SSID list on interfering APs (including known interfering APs). <p>If the valid SSID list is empty, it is ignored and clients can connect to all SSIDs on both valid APs and interfering APs (including known interfering APs). Not adding an SSID to the valid SSID list exposes that SSID to honeypot attacks.</p>

Table 102 *Valid SSIDs with Multi-Tenancy and Misconfigured AP Protection*

Multi-Tenancy Protection	Misconfigured AP Protection	Client Connections
Enabled	Enabled	<p>If there are entries in the valid SSID list:</p> <ul style="list-style-type: none"> • Clients can connect to valid SSIDs on valid APs. • Clients cannot connect to valid SSIDs on interfering APs (including known interfering APs). • Clients cannot connect to SSIDs not in the valid SSID list on valid APs. • Clients can connect to SSIDs not in the valid SSID list on interfering APs. <p>If the valid SSID list is empty, it is ignored and clients can connect to all SSIDs on both valid APs and interfering APs (including known interfering APs). Not adding an SSID to the valid SSID list exposes that SSID to honeypot attacks.</p>
Disabled	Enabled	<p>If there are entries in the valid SSID list:</p> <ul style="list-style-type: none"> • Clients can connect to valid SSIDs on valid APs. • Clients can connect to valid SSIDs on interfering APs (including known interfering APs). • Clients cannot connect to SSIDs not in the valid SSID list on valid APs. • Clients can connect to SSIDs not in the valid SSID list on interfering APs. <p>If the valid SSID list is empty, it is ignored and clients can connect to all SSIDs on both valid APs and interfering APs (including known interfering APs). When Multi-Tenancy protection is disabled, the network is susceptible to honeypot attacks.</p>

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.

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 [“Manual Blacklisting” on page 485](#) 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 [“Authentication Failure Blacklisting” on page 485](#) for more information.
- A denial of service 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 [“Attack Blacklisting” on page 485](#) 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 [Chapter 30, “External Services Interface”](#) for more information.



Note: The External Services Interface feature requires the Policy Enforcement Firewall Next Generation (PEFNG) license installed in the controller.

Manual Blacklisting

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 denial of service 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 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>
```

Authentication Failure Blacklisting

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.



Note: 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 [“Blacklist Duration” on page 486](#).

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>
```

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 denial of service (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 enable 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
```

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>
```


Dell PowerConnect 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) in the process of 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 8. With the introduction of LACP, this number remains the same. In essence, a port-channel group (LAG) is created either statically or dynamically via LACP. This chapter contains:

- [“Important Points to Remember” on page 489](#)
- [“Configuring LACP” on page 489](#)
- [“Best Practices” on page 491](#)
- [“Sample Configuration” on page 492](#)

Important Points to Remember

- 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 to 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 (i.e. VLAN membership, trunk status etc)
- 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.

Configuring LACP

Two LACP configured devices exchange LACPDU to form a 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 ArubaOS Command Line Reference Guide.

In the CLI

LACPDU 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 to 7.

```
lacp group <group_number> mode {active | passive}
```

- Active mode—the interface is in 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 state. The port in a passive mode responds to negotiations requests from other ports that are in an active state. Ports in passive state respond to LACP packets.



Note: A port in a passive state cannot set up a port channel (LAG group) with another port in a passive state.

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 time out value is long (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. Range is 1 to 65535 and default is 255.

4. View your LACP configuration.

The port uses the group number +1 as the “actor admin key”. By default, all the ports use the long timeout value (90 seconds).

```
(TechPubs)#show lacp 0 neighbor
Flags: S - Device is requesting Slow LACPDU
       F - Device is requesting fast LACPDU
       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 (that is, the partner device is different than the other ports) or the neighborship timesout or can not exchange LACPDU's with the partner, the port status is displayed as “DOWN” (see the following example).

```
(TechPubs)#show lacp 0 internal
Flags: S - Device is requesting Slow LACPDU
       F - Device is requesting fast LACPDU
       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 menus to enter the LACP values.

Network > Port

Port Port-Channel

Port Selection

0 2 4 6 8 10 12 14 16 18 20 22
1 3 5 7 9 11 13 15 17 19 21 23 24 25

Configure Selected Port 1/0

Enable Port

Enable 802.3af Power Over Ethernet

Make Port Trusted

Port Mode Access Trunk

Enter VLAN(s) VLAN ID 62

Trusted

VLAN Firewall Policy in

LACP

2 Group

active Mode

526 Priority

long Timeout

Apply

Commands [Hide Commands](#)

```
interface fastethernet 1/0
 lacp group 2 mode active
 lacp port-priority 526
 lacp timeout long
!
```

- LACP Group—The link aggregation group (LAG) number; 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; range is 1 to 65535 Default 255
- Timeout—Time out value for the LACP session; Long, the default, is 90 seconds; short is 3 seconds

Best Practices

- 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 it's actor admin key.
- By default, all ports use long timeout values (90 seconds)

- 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.

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
!
```


This chapter describes management access and tasks for a user-centric network and includes the following topics:

- “Certificate Authentication for WebUI Access” on page 493
- “Management Password Policy” on page 500
- “Managed RFprotect Sensors” on page 502
- “Managing Certificates” on page 503
- “Configuring SNMP” on page 507
- “Configuring Logging” on page 509
- “Guest Provisioning” on page 510
- “Managing Files on the Controller” on page 521
- “Setting the System Clock” on page 524

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).



Note: Each controller can support a maximum of ten management users.

Configuring Certificate Authentication for WebUI Access

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 503.
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>
```

Public Key Authentication for SSH Access

The controller allows public key authentication of users accessing the controller using SSH. (The default is for username/password authentication.) 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).

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 505](#).
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.



Note: Best practices 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>
```

Radius Server Authentication

Radius Server Username/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
```

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 [“Radius Server Username/Password Authentication” on page 495](#). 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.

RADIUS Server Authentication with Server-Derivation Rule



Note: 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.



Note: 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.

Configuring a Value-of 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 **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 244](#).

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 “[Management Password Policy](#)” on page 500.

[Figure 104](#) is an example of how to reset the password. The commands in bold type are what you enter.

Figure 104 *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 105](#) details an example reconfigure the enable mode password. Again, the command you enter appear in bold type.

Figure 105 *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
```

Setting an Administrator Session Timeout

You can configure the number of seconds after which an Administrator’s WebUI or CLI session times out.

Setting a CLI Session Timeout

To define a timeout interval for a CLI session, use the command:

```
login session timeout <value>
```

In the above command, <val> can be any number of minutes from 5 to 60, inclusive. You can also specify a timeout value of 0 to disable CLI session timeouts.

Setting a WebUI Session Timeout

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.

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. In the **Profiles List** window, expand the menu.
3. Select **Mgmt Password Policy**.
4. Configure the settings described in [Table 103](#).

Table 103 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.
Minimum number of Special characters	The minimum number of special characters required in a management user password. Range: 0-10 characters. By default, there is no requirement for special characters in a password, and the parameter has a default value of 0. See "Allowed Characters in a Management User Password" on page 501 for a list of allowed and disallowed special 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.

Table 103 Management Password Policy Settings

Parameter	Description
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 lockout the user upon crossing the "lock-out" threshold	The number of minutes a user who has exceeded the maximum number of failed password attempts is locked out of the network. After this period has passed, the lockout is cleared without administrator intervention. Range: 1 min to 1440 min (24 hrs). Default: 3.
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.

5. Click **Apply** to save your settings.

The table below describes the characters allowed in a management user password. The disallowed characters cannot be used by any management user password, even if the password policy is disabled.

Table 104 Allowed Characters in a Management User Password

Allowed Characters	Disallowed Characters
exclamation point: !	Parenthesis: ()
underscore: _	apostrophe: '
at symbol: @	semi-colon: ;
pound sign: #	dash: -
dollar sign: \$	equals sign: =
percent sign: %	slash: /
caret: ^	question mark: ?
ampersand: &	
star: *	
greater and less than symbols: < >	
curled braces: { }	
straight braces: []	
colon :	
period: .	
pipe:	
plus sign: +	
tilde: ~	
comma: ,	
accent mark: `	

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
```

Managed RFprotect Sensors

When an Dell controller is present in an Dell RFprotect system, an Dell AP that is acting as an RFprotect sensor can be configured and managed from the controller. As a Managed Sensor, the Dell AP is managed by the controller but sends collected security data about the wireless environment to an RFprotect Server.

A Managed Sensor is visible in ArubaOS from both the WebUI and CLI. From the Dell controller, you can perform the following management functions on a Managed Sensor:

- provision a Managed Sensor
- place a Managed Sensor into an AP group, and configure the AP group
- view Managed Sensor configuration and operation state
- change the mode of operation of the Managed Sensor between an AP/AM and an RFprotect sensor and vice versa
- reboot the Managed Sensor
- obtain certain software and hardware statistics on the Managed Sensor
- perform some debugging of the Managed Sensor

To change an Dell AP to a Managed RFprotect Sensor, configure the following in ArubaOS:

- In the radio profile for the AP, change the operating mode from ap-mode or am-mode to sensor-mode.
- In the AP system profile, enter the IP address of the RFprotect Server. If there is a backup RFprotect Server, enter the IP address of the backup.

The following sections describe how to configure these items using the ArubaOS WebUI or CLI.

Setting RFprotect Sensor Mode in the Radio Profile



Note: For a dual-radio AP, setting one radio in sensor mode causes both radios to act as RFprotect sensors. Changing the mode of a radio from AP or AM to sensor or from sensor to AP or AM causes the AP to reboot.

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 **RF Management**.
4. In the Profiles list, select the **802.11a** or **802.11g** radio profile.
5. For Mode, select **sensor mode**.
6. Click **Apply**.

RFprotect Managed Sensors are shown in the **Network > RFprotect Sensors and Controller > RFprotect Sensors** pages.

In the CLI

```
rf dot11a|dot11g-radio-profile <profile>
mode sensor-mode
```

In the outputs of the `show ap database` and `show ap active`, sensor mode is indicated with an “S” flag (for RFprotect Sensor).

Configuring the RFprotect Server

You can configure RFprotect server to monitor rogue APs and adhoc 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 **AP**, then **AP system profile**. The configuration setting are displayed in Profile Details.
4. Under Profile Details:
 - a. In the RFprotect Server IP field, enter the IP address of the server.
 - b. Optionally, in the RFprotect Backup Server IP field, enter the IP address of the backup RFprotect server.
 - c. Click **Apply**.

In the CLI

```
ap system-profile <profile>
rfprotect-server-ip <ipaddr>
rfprotect-bkup-server <ipaddr>
```

Reverting Managed Sensors to APs

To revert an Dell AP acting as a Managed RFprotect Sensor back to AP or AM mode, use the CLI or WebUI to change the operating mode of the AP in the radio profile (see [“Setting RFprotect Sensor Mode in the Radio Profile” on page 502](#)).

Managing Certificates

The Dell 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 (see [Chapter 9, “802.1x Authentication”](#)), VPN (see [Chapter 15, “Virtual Private Networks”](#)), 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).

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

Dell strongly recommends that you 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 505](#).



Note: 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. Click **Generate New**.
3. Enter the information detailed in [Table 105](#):

Table 105 *CSR Parameters*

Parameter	Description	Range
key	Length of private/public key.	1024/2048/4096
common_name	Typically, this is the host and domain name, as in <code>www.yourcompany.com</code> .	—
country	Two-letter ISO country code for the country in which your organization is located.	
state_or_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	Email address referenced in the CSR.	

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 key {1024|2048|4096} common-name <value> country <country>
state_or_province <state> city <city> organization <org> unit <string> email <email>
```

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

You must use the WebUI to import certificates into the controller. You cannot use a CLI command to import certificates, although a 'crypto-local pki' command is saved to the configuration file when you import a certificate from the WebUI.



Note: You cannot export certificates from the controller.

You can import the following types of certificates into the controller using the WebUI:

- 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 106 *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 107 *Imported Certificate Locations*

Location	Description
/flash/certmgr/trustedCAs	Trusted CA certificates, either for root or intermediate CAs. Dell recommends that if you 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.



Note: The controller does not support download of CRLs.

Configuring SNMP

Dell controllers support versions 1, 2c, and 3 of Simple Network Management Protocol (SNMP) for reporting purposes only.



Note: Dell-specific management information bases (MIBs) describe the objects that can be managed using SNMP. See the *Dell PowerConnect ArubaOS 5.0 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 108 *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.

Table 108 *SNMP Parameters for the Controller (Continued)*

Field	Description
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 Dell 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: <ul style="list-style-type: none"> • IP address • SNMP version: can be 1 or 2c. • Community 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 Dell 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: <ul style="list-style-type: none"> • MD5: HMAC-MD5-96 Digest Authentication Protocol • SHA: HMAC-SHA-96 Digest Authentication Protocol
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 Dell 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]
```


Configuring Logging

This section outlines the steps required to configure logging on an Dell controller. For each category or subcategory of message, you can set the logging level or severity level of the messages to be logged. [Table 109](#) summarizes these categories:

Table 109 *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
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 110](#) describes the logging levels in order of severity, from most to least severe.

Table 110 *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 **Add** 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>]
```

Guest Provisioning

The Guest Provisioning feature lets you manage guests who need access to your company's Dell 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



Note: 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 appears 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 106 Guest Provisioning Configuration Page—Guest Fields Tab

Management > Guest Provisioning

Guest Fields Page Design Email

Specify the fields you wish to appear on the Guest Provisioning Page. [Help](#)

Field	Internal Name	Label in UI	Display In	
			Details	Listing
	guest_category	Guest	<input checked="" type="checkbox"/>	<input type="checkbox"/>
	guest_username	Username	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	guest_password	Password	<input checked="" type="checkbox"/>	<input type="checkbox"/>
	guest_fullname	Full name	<input type="checkbox"/>	<input type="checkbox"/>
	guest_company	Company	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	guest_email	Email	<input checked="" type="checkbox"/>	<input type="checkbox"/>
	guest_phone	Phone	<input checked="" type="checkbox"/>	<input type="checkbox"/>
	comments	Comments	<input checked="" type="checkbox"/>	<input type="checkbox"/>
	account_category	Account	<input checked="" type="checkbox"/>	<input type="checkbox"/>
	creation_date	Created	<input checked="" type="checkbox"/>	<input type="checkbox"/>
	start_date	Start	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

2. Select the checkbox next to each field, described in [Table 111](#), that you want to appear on the Guest Provisioning page. Optionally, you can customize the label that appears 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: Dell recommends 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 111 Guest Provisioning—Guest Field Descriptions

Guest Field	Description
guest_category	A guest is the person who needs guest access to the company's Dell 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.

Table 111 Guest Provisioning—Guest Field Descriptions (Continued)

Guest Field	Description (Continued)
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	Username of the sponsor.
sponsor_dept	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 107 Guest Provisioning Configuration Page—Page Design Tab

Management > Guest Provisioning

Guest Fields | Page Design | Email

Specify the logo and colors for Guest Provisioning Page. [Help](#)

Banner:

Text:

Text Color:
 (RGB-6 Hex digits)

Background color:
 (RGB-6 Hex digits)

2. Enter the filename which contains the company banner in the **Banner** field. Or, click **Browse** to search for the filename



Note: Dell recommends using 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 appears 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 514](#)
 - [“Configuring an SMTP server and port in the CLI” on page 514](#)
2. Create the email messages. Complete this step using the WebUI:
[“Creating Email Messages in the WebUI” on page 514](#)

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 profile 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 Profile) #
(host) (Guest-access Email Profile) #smtp-port 25
(host) (Guest-access Email Profile) #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 108 Guest Provisioning Configuration Page—Email Tab

Management > Guest Provisioning

Guest Fields Page Design Email

Specify optional message to be sent to the guest and sponsor. [Help](#)

Guest Message

Subject: Guest account information

From: guest_admin@arubanetwork.com

Body: A guest account has been created for your use. The username, password and valid dates for the account are given below.

Send automatically at account creation time

Sponsor Message

Subject: Guest account information

From: sponsor_admin@arubanetworks.com

Body: You are listed as the Sponsor for the following guest account.

Send automatically at account creation time

2. To create a message for a guest or sponsor, customize the text in the **Subject**, **From** 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.



Note: 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 109 displays a sample email message that is sent to the guest after the guest account is created.

Figure 109 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



Note: 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: Dell recommends using 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 110](#).)

Figure 110 Customized Guest Account Information Window



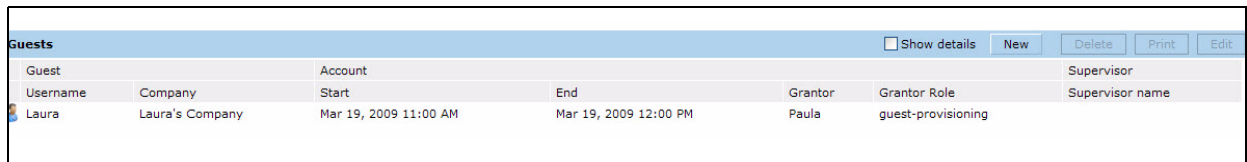
Creating Guest Accounts

After the Guest Provisioning user is created, that person can log in to the controller using the preconfigured username and password. The Management User Summary page displays. (See [Figure 112](#).) This is a sample page as the fields may differ based on how the network administrator designed the page.



Note: 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 111 *Creating a Guest Account—Management User Summary Page.*



Guests							<input type="checkbox"/> Show details	New	Delete	Print	Edit
Guest	Company	Account		End	Grantor	Grantor Role	Supervisor				
Username		Start					Supervisor name				
Laura	Laura's Company	Mar 19, 2009 11:00 AM		Mar 19, 2009 12:00 PM	Paula	guest-provisioning					



Note: If you do not want multiple guest users to share the same guest account concurrently, navigate to the Captive Portal Authentication profile 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

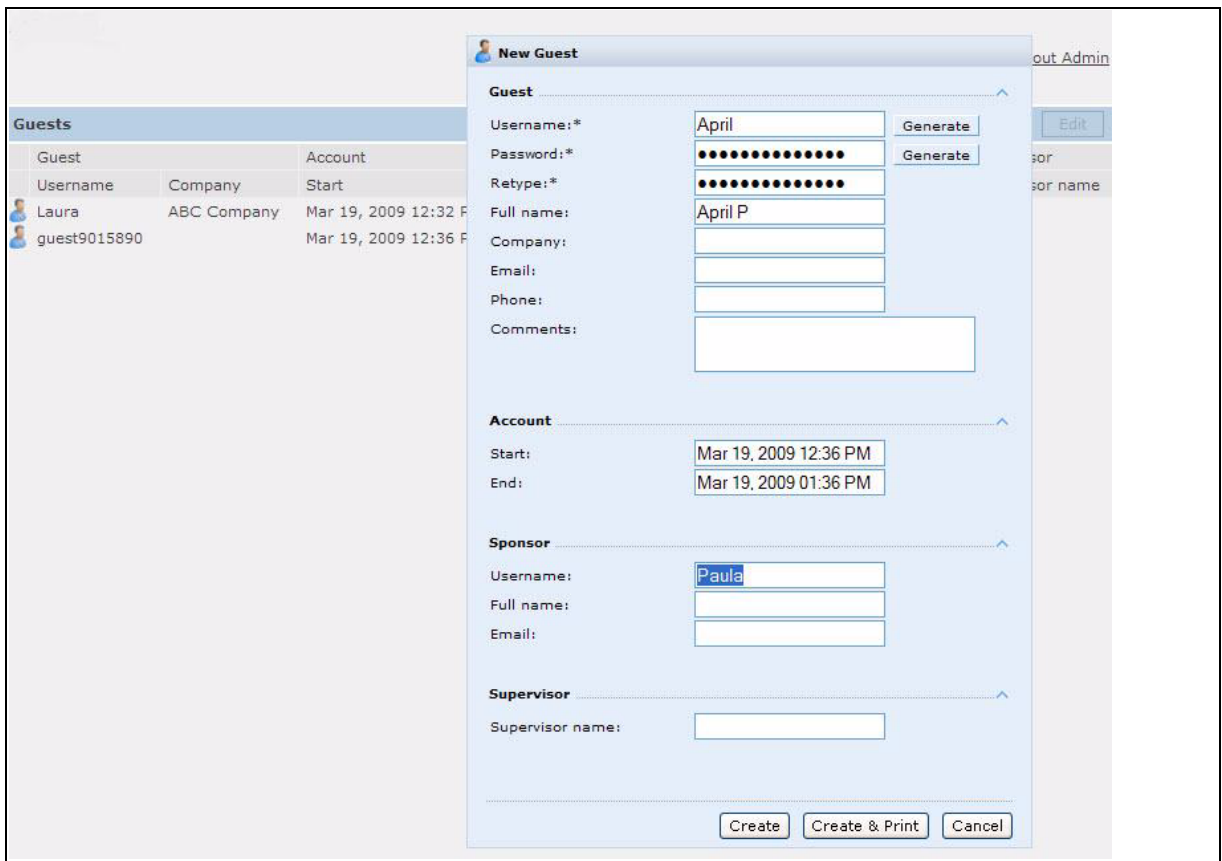
The Guest Provisioning user creates guest accounts by filling in information on the Guest Provisioning page. Tasks include creating, editing, enabling, printing, disabling and deleting guest accounts.

To create a new guest account, the Guest Provisioning user clicks **New** to display the New Guest window. (See [Figure 112](#).) After filling in information into the fields, click **Create**. The guest account now appears on the Management User Summary 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 112 *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. The Guest Provisioning user can send out Email from this window. (See [Figure 113](#).)

Figure 113 *Creating a Guest Account—Show Details Pop-up Window*



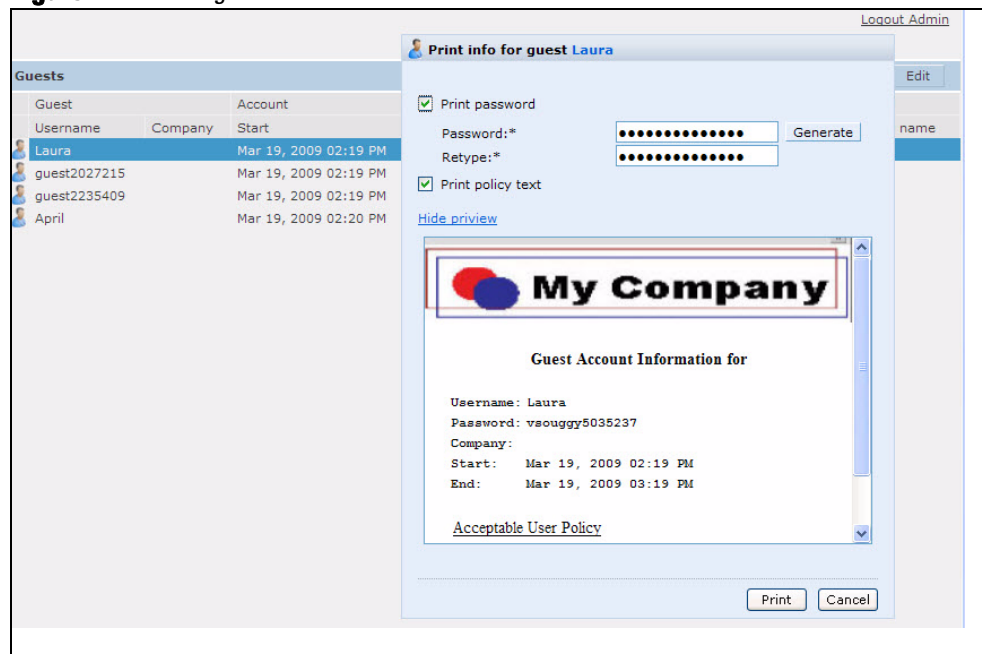
Printing Guest Account Information

To print guest account information:

1. Highlight the guest account you want to print and click **Print**. The **Print info** for guest window appears.

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 114](#).)
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 114 *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 Profiles** page.
2. Select **Wireless Lan**.
3. Under **Wireless Lan**, select and open **Captive Portal Authentication Profile**.
4. Add a new profile or select an existing profile.
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 <profile> 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.



Note: 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



Note: 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.



Note: 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.

[Table 112](#) lists the parameters that you configure to copy files to or from a controller.

Table 112 File Transfer Configuration Parameters

Server Type	Configuration
Trivial File Transfer Protocol (TFTP)	<ul style="list-style-type: none"> ● IP address of the server ● filename
File Transfer Protocol (FTP)	<ul style="list-style-type: none"> ● IP address of the server ● username and password to log into server ● filename
Secure Copy (SCP) You must use the CLI to transfer files with SCP.	<ul style="list-style-type: none"> ● IP address of the server or remote host ● username to log into server ● absolute path of filename (otherwise, SCP searches for the file relative to the user's home directory)

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 using 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> <password> <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>]
```

Configuring an NTP Server

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.



Note: 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]
```


Dell PowerConnect base features include sophisticated authentication and encryption, protection against rogue wireless As, 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, Policy Enforcement Firewall, AP Capacity. Evaluation licenses are available for some of these advanced features.

ArubaOS licenses are detailed in the following sections:

- “Terminology” on page 527
- “Licenses” on page 528
- “Multi-Controller Network” on page 529
- “License Usage” on page 529
- “Best Practices” on page 530
- “Installing a License” on page 531
- “Deleting a License” on page 532
- “Moving Licenses” on page 533
- “Resetting the Controller” on page 533

Terminology

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

- Bundles—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 (chassis or supervisor card) 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.

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 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.
- Policy Enforcement Firewall Virtual Private Network (PEFV)—Enables the roles based Policy Enforcement Firewall for VPN and 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, Voice and Video. This is an AP count license.
- Wireless Intrusion Protection (WIP)—Detection and prevention of wireless attacks, ad-hoc networks, signatures, denial of service attacks (DoS), impersonation, and misconfigured devices. This is an AP count license.
- Internal Test Functions—an internal license for internal use only.

License Types

These are the license categories available:

- **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.



Note: When license keys are applied on a controller, abnormal tampering of the device’s system clock (setting the system clock back by 2 hours or more) 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.

WLAN Network Status			
	Total Up	Total Down	IPSEC Up / IPSEC Down
WLAN Controllers	1	0	
Access Points	1	1659	0 / 0
Air Monitors	0	0	0 / 0
RFProtect Sensors	0	0	0 / 0
Wired Access Points	0	0	0 / 0
Disconnected Sensors	0		
Unprovisioned Access Points	0		
Duplicate AP Name	1000		

WLAN Attack Summary			
	Last 5 Min	Last Hour	All
Denial of Service Attacks	0	0	141
Impersonation Attacks	0	0	33
Signature Pattern Matches	0	0	21
Policy Violations	0	0	0
Unauthorized Devices Detected	0	0	741

Rogue AP Classification Summary			
	Last 5 Min	Last Hour	All
Rogue APs Detected	0	0	0

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.

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. If present, a backup master must be licensed with the same features as the Master. Backup controllers are “hot-standby”, that is, the backup controller processes AP, traffic, etc. while standing by in backup mode.

Dell PowerConnect best practices is to install the same set of feature licenses on every controller in your network.

License Usage

Licenses are platform independent and can be installed on any Dell controller. Installation of the feature license unlocks that feature’s functionality for the maximum capacity of the controller. [Table 113](#) list how licenses are consumed on the controllers.

Table 113 License Usage per License

License	Basis	What Consumes One License
PEFNG	AP	One operational AP
WIP	AP	One operational AP
AP	AP	One operational LAN-connected or mesh AP that is advertising at least one BSSID (virtual-AP) or RAP

The controller licenses are variable-capacity (see [Table 114](#)).



Note: In [Table 114](#), the Remote AP count is equal to the total AP count for each controller; the Campus AP count is 1/4 of the total AP count.

Table 114 MIPS Controller AP Capacity

Controller	Total AP Count	Campus APs	Remote APs
W-3200	128	32	128
W-3400	256	64	256
W-3600	512	128	512
W-620	32	8	32
W-650	64	16	64
W-651	64	16	64

Interaction

The various licenses do require some equality and other important interactions.

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



Note: It is not possible to designate specific APs for WIP/non-WIP operations.

- Mesh portals/mesh points, with no virtual-APs, do not consume a WIP license
- If a Mesh node is also configured for client service (advertises a BSSID for example), it consumes one AP license

Best Practices

- 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 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 alphanumeric 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. Or you can install licenses features yourself.



Note: Dell recommends that you obtain a user account on the Software License Management web site even if software license keys are pre-installed 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 software license from your sales account manager or authorized reseller (see [“Software License Email” on page 531](#)).
2. Locate the system serial number of your controller (see [“Locating the System Serial Number” on page 531](#)).

3. Use your system's serial number to obtain a software license key from the Software License Management web site (see “[Obtaining a Software License Key](#)” on page 532).
4. Enter the software license key via the controller's WebUI; 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 532).

Or

Launch the License Wizard from the **Configuration** tab and click the **New** button. Enter the software license key in the space provided (see “[Applying the Software License Key in the License Wizard](#)” on page 532).

5. Reboot your controller to enable your new license and features.

Software License Email

To obtain either a permanent or 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.



Note: 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 alphanumeric string used to access the license management Web site and which, in conjunction with the serial number of your controller, generates a unique software license key

Locating the System Serial Number

Each controller have unique serial numbers located at the rear of the controller . The location of the serial number is:

- at the rear of an Dell controller chassis

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.

Obtaining a Software License Key

To obtain a software license key, you must log in to the Dell License Management Web site at:

<https://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.

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

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



Note: The software license key is *only* valid 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: Log in to your controller's WebUI.

1. Navigate to the **Configuration > Network > Controller > System Settings** page and select the **License** tab.
2. Copy the software license key, from your email, and paste it into the **Add New License Key** field. Click **Add**.
3. Reboot your controller to enable the new license feature.

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 the **New** button.
2. The License Wizard will step you through the activation process. Click on the **Help** tab within the License Wizard for additional assistance.
3. Reboot your controller to enable the new license feature.

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 the **Delete** button at the far right hand side of the license to delete the license.

If a license feature is under an evaluation license, no key is generated when the feature is deleted.

Moving Licenses

It may become necessary to move licenses from one controller to another or simply delete the license for future use. To move licenses, delete the license from the chassis as described in [“Deleting a License” on page 532](#). Then install the license key on the new controller as described in [“Applying the Software License Key in the WebUI” on page 532](#).



Note: The ability to move a license from one controller to another is provided to allow customers maximum flexibility in managing their organization's network and to minimize the need to contact Dell customer support. License fraud detection is monitored and enforced by Dell PowerConnect. Abnormally high volumes of license transfers for the same license certificate to multiple controllers can indicate 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 permanent or evaluation licenses.

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



Warning: 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.

This chapter describes ArubaOS support for IPv6 clients.

- [“About IPv6” on page 535](#)
- [“Support for IPv6” on page 535](#)
- [“Features that Support IPv6” on page 537](#)
- [“IPv6 User Addresses” on page 542](#)
- [“Important Points to Remember” on page 542](#)

About IPv6

The IPv6 protocol enables the next generation of large-scale IP networks by supporting addresses that are 128 bits long. This allows for 2^{128} possible addresses (versus 2^{32} possible IPv4 addresses).

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

```
FEDC:BA98:7654:3210:FEDC:BA98:7654:3210  
1080:0:0:0:0:800:200C:417A
```

The use of the “::” symbol is a special syntax that you can use to compress one or more 16-bit groups of zeros or to compress leading or trailing zeros in an address. The “::” can appear only once in an address. For example, the following example address:

```
1070:0:0:0:0:800:200C:417A
```

can also be represented as:

```
1080::800:200C:417A
```

IPv6 uses subnet identifiers to identify subnetworks to which nodes are attached. In ArubaOS, when you reference IPv6 subnetworks in firewall policies, you must specify a subnet mask in addition to the IPv6 address. The subnet mask is a bitmask that specifies the prefix length. For example, the IPv6 address and subnet mask:

```
1080::800:200C:417A ffff:ffff:ffff:ffff::
```

represents all IPv6 addresses with the subnet identifier 1080:0:0:0.

Support for IPv6

ArubaOS provides wired or wireless clients using IPv6 addressing 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 [“Important Points to Remember” on page 542](#)).

Supported Network Configuration


Clients can be wired or wireless and use IPv4 and/or IPv6 addressing. ArubaOS requires that the default gateway for the IPv6 clients be an external router that supports IPv6. The Dell controller itself has an IPv4 address, and cannot route packets with IPv6 addresses. You can use the WebUI or CLI to display IPv6 client information.

IPv6 clients must be mapped to a VLAN that is bridged to an external router which provides IPv6 services to those clients. On the controller, you can configure IPv4 and IPv6 clients on the same VLAN.



Note: IPv6 clients and the IPv6 router must be on the same VLAN.

Figure 115 *Supported Network Configuration*



Network Connection 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 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. Entries are created in the user and session tables.
3. 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. IPv6 client sends a Neighbor Solicitation message.
6. IPv6 router responds with a Neighbor Advertisement message.
7. If DHCP is required to provide IPv6 addresses, the DHCPv6 process is started.
8. 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. Entries are created in the user and session tables.

Features that Support IPv6

This section describes ArubaOS features that support IPv6 clients.

Authentication

ArubaOS only supports 802.1x authentication for IPv6 clients. You cannot configure layer-3 authentications such as captive portal to authenticate IPv6 clients.

Table 115 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	No
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 configuration. For more information about configuring 802.1x authentication on the controller, see [Chapter 9, “802.1x Authentication”](#) on page 251.



Note: This release does not support authentication of management users on IPv6 clients.

Firewall Functions

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 [“Global Firewall Parameters”](#) on page 289.



Note: Voice-related and NAT firewall functions are not supported for IPv6 traffic.

Table 116 IPv6 Firewall Parameters

Authentication Method	Description
Monitor Ping Attack	Number of ICMP pings per second, which if exceeded, can indicate a denial of service attack. Valid range is 1–255 pings per second. Recommended value is 4. Default: No default
Monitor TCP SYN Attack rate	Number of TCP SYN messages per second, which if exceeded, can indicate a denial of service attack. Valid range is 1–255 messages per second. Recommended value is 32. Default: No default
Monitor IP Session Attack	Number of TCP or UDP connection requests per second, which if exceeded, can indicate a denial of service attack. Valid range is 1–255 requests per second. Recommended value is 32. Default: No default

Table 116 IPv6 Firewall Parameters (Continued)

Authentication Method	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 All IP Fragments	Drops all IP fragments. NOTE: Do not enable this option unless instructed to do so by an 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, 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 NOTE: An IPv6 client can have multiple IP addresses. Enabling IP spoofing on the controller can cause IPv6 clients to lose network access.
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 an 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 an 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 an Dell representative, as doing so may create unnecessary overhead on the controller. Default: Disabled (per-session logging is performed)

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
```

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 one or more 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 for IPv6 traffic and apply IPv4 and IPv6 firewall policies to the same user role. For example, if you have employees that are using 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 117](#) describes required and optional parameters for an IPv6 firewall policy rule.

Table 117 IPv6 Firewall Policy Rule Parameters

Field	Description
Source (required)	<p>Source of the traffic, which can be one of the following:</p> <ul style="list-style-type: none"> <i>any</i>: Acts as a wildcard and applies to any source address. <i>user</i>: This refers to traffic from the wireless client. <i>host</i>: 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. <i>network</i>: 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 IPv6 address and network mask of the subnet. For example, 2002:ac10:fe::ffff:ffff:ffff::. <i>alias</i>: This refers to using an alias for a host or network. <p>NOTE: This release does not support IPv6 aliases. You cannot configure an alias for an IPv6 host or network.</p>
Destination (required)	<p>Destination of the traffic, which can be configured in the same manner as Source.</p>
Service (required)	<p>NOTE: Voice over IP services are not available for IPv6 policies.</p> <p>Type of traffic, which can be one of the following:</p> <ul style="list-style-type: none"> <i>any</i>: This option specifies that this rule applies to any type of traffic. <i>tcp</i>: Using this option, you configure a range of TCP port(s) to match for the rule to be applied. <i>udp</i>: Using this option, you configure a range of UDP port(s) to match for the rule to be applied. <i>service</i>: 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. <i>protocol</i>: Using this option, you specify a different layer 4 protocol (other than TCP/UDP) by configuring the IP protocol value.
Action (required)	<p>The action that you want the controller to perform on a packet that matches the specified criteria. This can be one of the following:</p> <p>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).</p> <ul style="list-style-type: none"> <i>permit</i>: Permits traffic matching this rule. <i>drop</i>: Drops packets matching this rule without any notification.
Log (optional)	<p>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.</p>

Table 117 IPv6 Firewall Policy Rule Parameters (Continued)

Field	Description
Mirror (optional)	Mirrors session packets to datapath or remote destination specified in the IPv6 firewall function (see “Session Mirror Destination” in Table 116 on page 537). 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”.



Note: 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 **IPv6 Session** for Policy Type.
5. Click **Add** to add a rule that allows HTTP traffic.
 - a. Under Source, select **network** from the drop-down list.
 - b. For Host IP, enter **2002:d81f:f9f0:1000::**.
 - c. For Mask, enter **ffff:ffff:ffff:ffff::**.
 - d. Under Service, select **service** from the drop-down list.
 - e. Select **svc-http** from the scrolling list.
 - f. Click **Add**.
6. Click **Add** to add a rule that allows HTTPS traffic.
 - a. Under Source, select **network** from the drop-down list.
 - b. For Host IP, enter **2002:d81f:f9f0:1000::**.
 - c. For Mask, enter **ffff:ffff:ffff:ffff::**.
 - d. Under Service, select **service** from the drop-down list.
 - e. Select **svc-https** from the scrolling list.
 - f. Click **Add**.



Note: Rules can be reordered using the up and down arrow buttons provided for each rule.

7. Click **Apply** to apply the configuration. 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:

```
ipv6 access-list session ipv6-web-only
  network 2002:d81f:f9f0:1000:: ffff:ffff:ffff:ffff:: any svc-http permit
  network 2002:d81f:f9f0:1000:: ffff:ffff:ffff:ffff:: 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 apply this configuration.

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
```

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.

Multicast Snooping

Multicast Listener Discovery (MLD) protocol enables an IPv6 router to discover the presence of multicast listeners on directly-attached links. This release of ArubaOS supports version 1 of the MLD protocol (MLDv1). MLDv1, defined in RFC 2710, is derived from version 2 of the IPv4 Internet Group Management Protocol (IGMPv2). You can optionally enable MLD snooping to limit the sending of multicast frames to only those nodes that need to receive them. Protocol Independent Multicast (PIM) is not supported.

The following examples create VLAN 22 and enable MLDv1 and MLD snooping on the VLAN.

Use the procedure described below to enable MLDv1 via the WebUI.

1. Navigate to the **Configuration > Network > VLANs** page.
2. Click **Add** to create a new VLAN.
3. On the Add New VLAN page, enter 22 for the VLAN ID.
4. Click **Apply**.
5. Navigate to the **Configuration > Network > IP > IP Interfaces** page.
6. Click **Edit** for VLAN 22.
7. Select both **Enable MLD** and **Snooping**.

8. Click **Apply**.

To enable MLDv1 via the command-line interface, issue the following commands in config mode:

```
vlan 22
interface vlan 22
    ipv6 mld snooping
```

IPv6 User Addresses

Viewing or Deleting User Entries

There is a separate user table for IPv6 users that contains entries for every IPv6 address used by a client. 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, issue the **show ipv6 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
```

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 ipv6 datapath session**. To display the user entries in the datapath, access the command-line interface in enable mode, and issue the command **show ipv6 datapath user**. For details on each of these commands and the output they display, refer to the *Dell PowerConnect ArubaOS CLI Reference Guide*.

Important Points to Remember

This ArubaOS release does not support the following functions for IPv6 clients:

- Do not use VLAN pooling if you enable IPv6 forwarding on the controller, as VLAN pooling will flood IPv6 multicast packets for all VLANs that are part of the VLAN pool. This can cause autoconfigured clients to acquire multiple IPv6 addresses (one for each vlan in the pool) making those clients behave unpredictably. If you need to work around this limitation, you can unicast BC/MC traffic to every station. To enable this workaround, you must enable the `wlan ssid-profile battery-boost` option, and install a Policy Enforcement Firewall Next Generation (PEFNG) license.
- The controller cannot route packets with IPv6 addresses; the routing function must be performed by an external IPv6 router.
- The controller does not perform network address translation on IPv6 addresses.
- The controller does not generate any IPv6 ICMP messages.
- Voice over IP 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.
- The controller cannot terminate VPNs for IPv6 clients.
- Layer-3 authentications, such as captive portal and VPN authentication, cannot be performed for IPv6 clients.
- ArubaOS does not support RADIUS over IPv6 as an authentication protocol.

- Authentication of management users on IPv6 clients is not supported.
- The controller does not access the flow information field in IPv6 packet headers. (This field can be used by IPv6 routers to identify the sequence of packets and to facilitate routing decisions.)
- A client can have an both IPv4 address and an IPv6 address, but the controller does not relate the states of the IPv4 and IPv6 addresses on the same client. For example, if an IPv6 user session is active on a client, an IPv4 user session on the same client will be deleted if the idle timeout for the IPv4 session is reached.

This chapter outlines the steps required to configure voice and video services on Dell controller for Voice over IP (VoIP) devices, including Session Initiation Protocol (SIP), Spectralink Voice Priority (SVP), H323, SCCP, Vocera, and Alcatel NOE phones. Since 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 describes the following topics:

- “License Requirements” on page 545
- “Configuring Voice” on page 545
- “Configuring Video” on page 557

License Requirements

See [Chapter 27, “Software Licenses” on page 527](#) for more information on license requirements.

Configuring Voice

This section describes the steps required to set up and configure voice features on Dell controller. To configure voice features you must do the following:

1. Set up net services
2. Configure roles
3. Configure ALG
4. Configure other parameters depending on the need and environment



Note: Assigning voice traffic to the high priority queue is recommended when deploying voice over WLAN networks.

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 118 *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	UDP	5060	SIP
svc-sips	TCP	5061	SIP

Table 118 *Default Voice Net Services and Ports*

Net Service Name	Protocol	Port	ALG
svc-noe	UDP	32512	NOE
svc-h323-udp	UDP	1718, 1719	H.323
svc-h323-tcp	TCP	1720	H.323
svc-vocera	UDP	5002	VOCERA
svc-svp	SVP	None	SVP

Creating or Modifying 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 a *svc-noe-oxo* service on UDP port 5000, enter:

```
(host) (config)# netservice svc-noe-oxo udp 5000 alg noe
```

Configuring User Roles

In the Dell 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 [Chapter 10, “Roles and Policies,”](#) 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 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 **IPv4 Session**.
5. Under Rules, click **Add**.
 - a. For Source, select **any**.
 - b. For Destination, select **any**.
 - c. For Service, select service, then select the correct ALG service. See [Table 119 on page 547](#) for service names for all ALGs:

Table 119 Services for ALGs

ALG	Service Name
NOE	svc-noe
SIP	Select one of these service: <ul style="list-style-type: none"> • svc-sip-tcp • svc-sip-udp • svc-tftp
SVP	Select one of these service: <ul style="list-style-type: none"> • svc-scp • svc-tftp • svc-dhcp
VOCERA	svc-vocera
SCCP	Select one of these service: <ul style="list-style-type: none"> • svc-sccp • svc-tftp
H.323	Select one of these service: <ul style="list-style-type: none"> • svc-h323-tcp • svc-h323-udp

- 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 (step 3) 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 CLI to configure a user role

```
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 118 on page 545](#).

Using 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.



Note: 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

```
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

```
aaa derivation-rules user name
    set role condition macaddr contains xx:xx:xx set-value role
```

Optional Configurations

This section describes other voice-related features that you can configure in the base ArubaOS.

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.



Note: This feature supports WPA clients, while opportunistic key caching (also configured in the 802.1x Authentication profile) supports WPA2 clients.

Using the WebUI to enable WPA fast handover

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 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 to select the **WPA-Fast-Handover** check box.
5. Click **Apply**.

Using the CLI to enable WPA fast handover

```
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), Dell recommends that you enable the “local probe response” option in the SSID profile. (Generating probe responses on the Dell controller is an optimization that allows ArubaOS to make better decisions.) This option is enabled by default in the SSID profile. You can also 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.

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 [Chapter 19, “IP Mobility”](#) for more information about mobility.

The 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.

Using the WebUI to configure a VoIP Call Admission Control profile

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 120](#) describes the parameters you can configure in this profile

Table 120 VoIP Call Admission Control Configuration Parameters

Parameter	Description
VoIP Call Admission Control	Select the Voip Call Admission Control checkbox to enable WiFi 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. The default value is 10. 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 .
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. The default value is 2000 kbps.
VoIP Call Handoff Reservation	Specify the percentage of call capacity reserved for mobile VoIP clients on an active call. The default value is 20%.

Table 120 VoIP Call Admission Control Configuration Parameters

Parameter	Description
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. Select the VoIP Send SIP 100 Trying checkbox to send <i>SIP 100-trying</i> 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 the disconnect extra call feature is enabled, 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 number of seconds that a client must wait to start the call after sending the TSPEC request (the default is one second). Select the VoIP TSPEC Enforcement checkbox to validate of 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: <ul style="list-style-type: none"> ● 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: <ul style="list-style-type: none"> ● 480: Temporary Unavailable ● 486: Busy Here ● 503: Service Unavailable ● none: Don't send SIP status code

5. Click **Apply** to save your settings.

Using the CLI to configure the VoIP Call Admission Control profile

```
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>
```

VoIP-Aware ARM Scanning

The VoIP-aware ARM scanning feature allows you to provide higher QoS to the voice traffic. You can use the WebUI or CLI to enable VoIP-aware ARM scanning.

Using the WebUI to enable VoIP aware scanning in the ARM profile

1. Navigate to the **Configuration > AP Configuration** page. Select either the **AP Group** or **AP Specific** tab.
 - If you selected the **AP Group** tab, click the **Edit** button by the name of the AP group with the ARM profile you want to configure.
 - If you selected the **AP Specific** tab, click the **Edit** button 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 (check) the **VoIP Aware Scan** option.
6. Click **Apply**.

For additional information on configuring an Adaptive Radio Management profile, see [“ARM Profiles” on page 140](#).

Using the CLI to enable VoIP aware scanning in the ARM profile

```
rf arm-profile <profile-name>  
  voip-aware-scan
```

Voice-Aware 802.1x

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.

Using the WebUI to disable voice awareness for 802.1x

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**.

Using the CLI to disable voice awareness for 802.1x

```
aaa authentication dot1x <profile>  
  no voice-aware
```

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 (the default role is guest). You specify a configured user role for the SIP client in the AAA profile.

Using the WebUI to configure the SIP client user role

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 select 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**.

Using the CLI to configure the SIP client user role

```
aaa profile <profile>
  sip-authentication-role <role>
```

Use the `show voice sip client-status` command to view the state of the client registration.

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.



Note: Dial plan can be configured only for SIP over UDP.

Dial Plan Format

The format of a SIP dial plan is `<sequence> <digit-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.
- `digit-pattern`—is the 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 dial plan that is automatically prefixed to the dialed number. This is specified as `<dial-plan>%e`. Examples of dial plans are:
 - 9%e: The number 9 is prefixed to the dialed number.
 - 91%e: The number 91 is prefixed to the dialed number.

Table 121 *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	his 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.

Using CLI to configure dial plan

To configure a dial plan for SIP devices:

1. Create a voice dial plan
2. Associate the dial plan with SIP ALG

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") #!
```

Associate the dial plan with SIP ALG

```
(host) (config) #voice sip
(host) (SIP settings) #dialplan-profile local
(host) (SIP settings) #!
```

View SIP dial plan profile


```
(host) (config) #show voice sip
SIP settings
-----
Parameter      Value
-----
Dialplan Profile local
```

View dial plan details


```
(host) (config) #show voice dialplan-profile local
Dialplan Profile "local"
-----
Parameter Value
-----
dialplan  100 XXXXXXX 9%e
```

Using WebUI to configure dial plan

1. In the WebUI, navigate to Configuration > Advanced Services > All Profiles > Controller > Dialplan Profile. Enter a name for the dial plan profile and click the Add button.

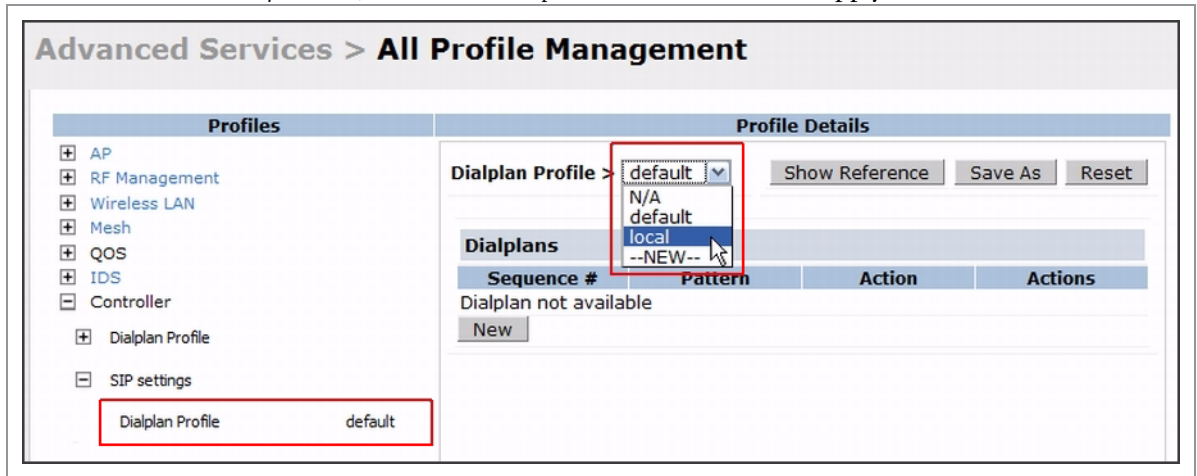


2. Under *Profiles*, expand **Controller** and select the newly created dial plan profile. Enter the following dial plan details and click the **Add** button.
 - Sequence number: The dial plan position in the list of dial plans.
 - Pattern: The number that the user will dial.
 - Action: Prefix to be added by the controller before forwarding the call to the EPABX.

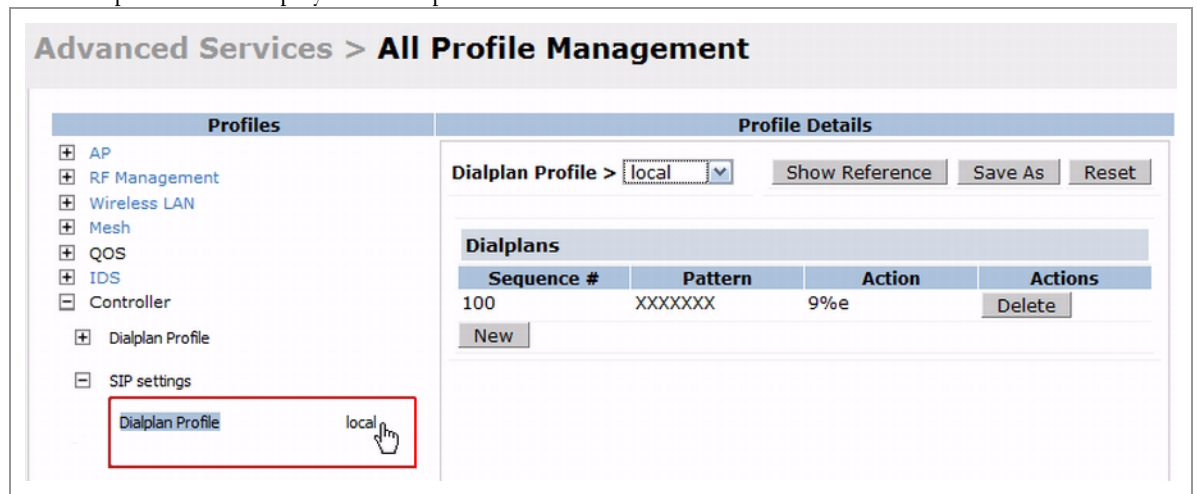


Click the **Apply** button to save the configuration.

- 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 the **Apply** button.



The Dialplan Profile displays the dial plan details:



Voice over Remote Access Point

Voice traffic support is enhanced on split tunnel mode over 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 now work reliably in split tunnel mode
- 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		
moscatp	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
```

Configuring Video

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 loss. 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, do the following:

- Enable WMM on the SSID profile.
- Enable IGMP proxy settings.
- Set a DSCP value for the video stream—All streams with DSCP value are sent for conversion.
- Create an ACL for the multicast group with the DSCP value that is same as the `wmm-vi-dscp` value in SSID profile.
- Configure dynamic multicast optimization—All streams with the DSCP value sent to an AP are dynamically optimized for streaming.
- 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.
- Configure ARM scanning for video traffic—This ensures that AP does not scan when a video stream is active.
- Optionally you can configure and apply WMM bandwidth management profile—The total bandwidth share should not exceed 100 percent.

You can either use CLI or WebUI to configure the video over WLAN enhancements.

Pre-requisites

- You will need the Policy Enforcement Firewall Next Generation (PEFNG) license to enable dynamic multicast optimization.
- This feature is available only on W-3000 Controller Series, W-600 Series Controller, and W-3400 controller platforms.

Using CLI

1. Set a DSCP value for video traffic.

```
(host) (config)#wlan ssid-profile default
(host) (ssid-profile "default" )#wmm-vi-dscp <value>
```

Example: (host) (ssid-profile "default")#wmm-vi-dscp 40

Setting the DSCP value, tags the content as video stream that the APs can recognize. By default, the DSCP value is set to 40. You must also set an ACL on the controller with equivalent 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-video)#any network 224.0.0.0 255.0.0.0 any permit tos 40 queue high
```

2. 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
...
...
```

3. Configure the dynamic multicast optimization threshold value.

```
(host) (config) #dynamic-mcast-optimization-thresh 6
(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
...
...
```

4. 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  
Max Tx EIRP                              127 dBm  
Min Tx EIRP                              9 dBm  
Multi Band Scan                          Enabled  
Rogue AP Aware                           Disabled  
Scan Interval                             10 sec  
Active Scan                               Disabled  
Scanning                                  Enabled  
Scan Time                                 110 msec  
VoIP Aware Scan                           Disabled  
Power Save Aware Scan                     Enabled  
Video Aware Scan                          Enabled  
Ideal Coverage Index                      10  
Acceptable Coverage Index                 4  
Free Channel Index                        25  
Backoff Time                              240 sec  
Error Rate Threshold                      50 %  
Error Rate Wait Time                      30 sec  
Noise Threshold                           75 -dBm  
Noise Wait Time                           120 sec  
Minimum Scan Time                          8  
Load aware Scan Threshold                 1250000 Bps  
Mode Aware Arm                            Disabled
```

5. Configure and apply a bandwidth management profile.

```
(host) (config)# #wlan wmm-traffic-management-profile default
```



Note: Configure the virtual AP traffic management profile before applying the WMM traffic management profile to the virtual AP profile.

- a. Enable a bandwidth shaping policy so that the allocated bandwidth share is appropriately used.

```
(host) (WMM Traffic management profile "default") # enable-shapping
```

- 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
```

Using WebUI

To access the WebUI configuration screens navigate to **Configuration > Advanced Services > All Profiles**.

1. Set a DSCP value for video traffic.

Under the **Profiles** column, expand **Wireless LAN > SSID Profile** and select the profile name. This example uses the *default* profile. Enter the DSCP value (integer number) and click the **Apply** button.

Figure 116 Setting DSCP value

Profile Details	
Max Transmit Attempts	8
Short Preamble	<input checked="" type="checkbox"/>
Wireless Multimedia (WMM)	<input type="checkbox"/>
WMM TSPEC Min Inactivity Interval	0 msec
DSCP mapping for WMM video AC	32
DSCP mapping for WMM background AC	8
Hide SSID	<input type="checkbox"/>
Local Probe Response	<input checked="" type="checkbox"/>
Battery Boost	<input type="checkbox"/>
WEP Key 2	Retype: <input type="text"/>
WEP Key 4	Retype: <input type="text"/>
RTS Threshold	2333 bytes
Max Associations	64
Wireless Multimedia U-APSD (WMM-UAPSD) Powersave	<input checked="" type="checkbox"/>
DSCP mapping for WMM voice AC	56
DSCP mapping for WMM best-effort AC	24
902# Competibility Mode	<input type="checkbox"/>
Deny_Broadcast Probes	<input type="checkbox"/>
Disable Probe Retry	<input checked="" type="checkbox"/>
WEP Key 1	Retype: <input type="text"/>
WEP Key 3	Retype: <input type="text"/>
WEP Transmit Key Index	1

2. 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 117 Enabling Dynamic Multicast Optimization for Video and Set Threshold

Virtual AP profile > default	
Virtual AP enable	<input checked="" type="checkbox"/>
VLAN	<-- --NONE-- >
Deny time range	--NONE--
HA Discovery on-association	<input type="checkbox"/>
Station Blacklisting	<input checked="" type="checkbox"/>
Dynamic Multicast Optimization (DMO)	<input checked="" type="checkbox"/>
Authentication Failure Blacklist Time	3600 sec
Strict Compliance	<input type="checkbox"/>
Remote-AP Operation	standard
Convert Broadcast ARP requests to unicast	<input type="checkbox"/>
Allowed band	all
Forward mode	decrypt-tunnel
Mobile IP	<input checked="" type="checkbox"/>
DoS Prevention	<input type="checkbox"/>
Blacklist Time	3600 sec
Dynamic Multicast Optimization (DMO) Threshold	6
Multi Association	<input type="checkbox"/>
VLAN Mobility	<input type="checkbox"/>
Drop Broadcast and Multicast	<input type="checkbox"/>
Band Steering	<input type="checkbox"/>

3. 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 the **Apply** button.

Figure 118 Enabling Video Aware Scan



4. Configure and apply bandwidth management profile

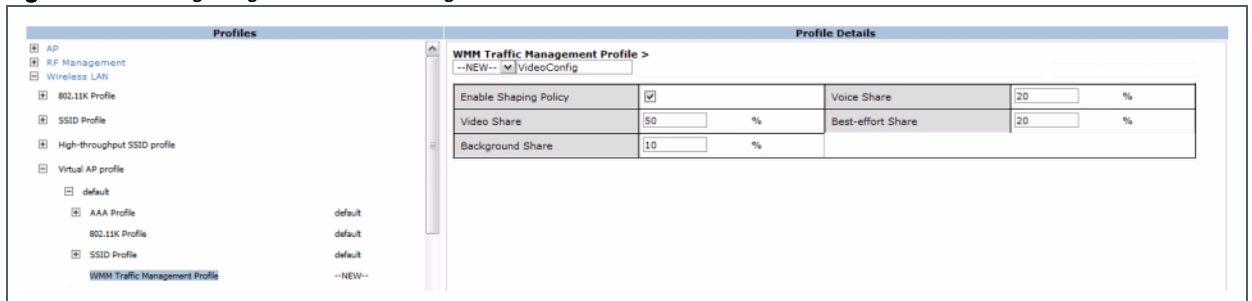
Under the **Profiles** column, expand **Virtual AP > [profile-name] > WMM Traffic Management Profile**. In the **Profile Details** section, select the profile name from the drop down list box. Select the **Enable Shaping Policy** option and enter the bandwidth share values. Click the **Apply** button to save the settings.

This step is optional.



Note: Ensure that you configure the virtual AP traffic management profile before applying the WMM traffic management profile to the virtual AP profile.

Figure 119 Configuring bandwidth management



After you configure the WMM bandwidth management profile, apply it to the virtual AP profile.

QoS

Thus, QoS for voice and video applications is configured when you configure firewall roles and policies

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.



Note: WMM does not support APs configured in bridge mode.

WMM supports four access categories (ACs): voice, video, best effort, and background. <_blue> on page 562 shows the mapping of the WMM access categories to 802.1D priority values. The 802.1D priority value is contained in a two-byte QoS control field in the WMM data frame.

Table 122 WMM Access Category to 802.1D Priority Mapping

Priority	802.1D Priority	WMM Access Category
Lowest ↓ Highest	1	Background
	2	
	0	Best effort
	3	
	4	Video
	5	
	6	Voice
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 those environments in which the wireless clients support WMM, you can enable both WMM and U-APSD in the SSID profile.

Using the WebUI to enable WMM

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**. Select **Virtual AP**, then select 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 (check) this option.
6. Click **Apply**.

Using the CLI to enable WMM

```
wlan ssid-profile <profile> wmm  
wlan ssid-profile <profile> wmm-uapsd
```

Configurable 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.



Note: The user-configured mapping only takes effect when WMM is enabled for the SSID profile.

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 123 on page 563](#) shows the default WMM AC to DSCP decimal mappings and the recommended WMM AC to DSCP Hex mappings.

Table 123 WMM Access Category to DSCP Mappings

DSCP Decimal Value (default mappings)	DSCP Hex Value (recommended mappings)	WMM Access Category
8	0x08	Background
	0x10	
24	0x00	Best effort
	0x18	
40	0x20	Video
	0x28	
56	0x30	Voice
	0x38	

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.



Note: 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.

Mapping Considerations

If you do not define a mapping for a particular DSCP tagged packet, default mappings are applied and prioritized accordingly (Best Effort uses 0x00).



Note: The WMM AC mapping commands do not take affect on APs configured in bridge mode.

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**. Select **Virtual AP**, then select 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 (check) 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**.

Using the CLI to map between WMM AC and DSCP

```
wlan ssid-profile <profile>
  wmm-be-dscp <best-effort>
  wmm-bk-dscp <background>
  wmm-vi-dscp <video>
  wmm-vo-dscp <voice>
wmm
```

Battery Boost

Battery boost is an optional feature that can be enabled for any SSIDs that support voice traffic. This feature converts all multicast traffic to unicast before delivery to the client. Enabling battery boost on an SSID allows you to set the DTIM interval from 10 - 100 (the previous allowed values were 1 or 2), equating to 1,000 - 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.

The following step enables the battery boost feature and sets the DTIM interval in the SSID profile.

Using the WebUI to enable battery boost

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**, expand **Wireless LAN**, then select **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**.

Using the CLI to enable battery boost

```
wlan ssid-profile <profile>
  battery-boost
  dtim-period <milliseconds>
```


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. WMM anticipates the ratification of the IEEE 802.11e standard that is currently in development.

WMM requires:

- The access point is Wi-Fi Certified and has WMM enabled
- The client device is Wi-Fi Certified
- The application supports 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 802.1d priority tags, as shown in [Table 124 on page 565](#).

Table 124 WMM Access Categories and 802.1d Tags

WMM Access Category	Description	802.1d 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 as 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 affect traffic from the AP to the client.
- STA parameters affect 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**, expand the **Wireless LAN** menu, then select **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 125](#) describes the parameters you can configure in this profile.

Table 125 EDCA Parameters Station and EDCA Parameters AP Profile Settings

Parameter	Description
Best Effort	<p>Set the following parameters to define the best effort queue.</p> <ul style="list-style-type: none"> • aifsn: Arbitrary inter-frame space number. Possible values are 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 $2^4-1 = 15$. Possible values are 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 $2^4-1 = 15$. Possible values are 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). Possible values are 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	<p>Set the following parameters to define the background queue.</p> <ul style="list-style-type: none"> • aifsn: Arbitrary inter-frame space number. Possible values are 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 $2^4-1 = 15$. Possible values are 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 $2^4-1 = 15$. Possible values are 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). Possible values are 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	<p>Set the following parameters to define the background queue.</p> <ul style="list-style-type: none"> • aifsn: Arbitrary inter-frame space number. Possible values are 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 $2^4-1 = 15$. Possible values are 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 $2^4-1 = 15$. Possible values are 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). Possible values are 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.

Table 125 EDCA Parameters Station and EDCA Parameters AP Profile Settings

Parameter	Description
Voice	<p>Set the following parameters to define the background queue.</p> <ul style="list-style-type: none">● aifsn: Arbitrary inter-frame space number. Possible values are 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 $2^4-1 = 15$. Possible values are 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 $2^4-1 = 15$. Possible values are 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 ($3008/32$). Possible values are 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

```
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>
```

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 this feature is enabled, 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 is used to inform the client that the reservation is terminated.

Using the WebUI to enable WMM queue content enforcement

1. Navigate to the Configuration > Advanced Services > Stateful Firewall page.
2. Select Enforce WMM Voice Priority Matches Flow Content.
3. Click Apply.

Using the CLI to enable WMM queue content enforcement

```
firewall wmm-voip-content-enforcement
```


The Dell 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 Dell 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.

This chapter describes the following topics:

- [“Understanding ESI” on page 569](#)
- [“Understanding the ESI Syslog Parser” on page 571](#)
- [“ESI Configuration Overview” on page 574](#)
- [“Example Route-mode ESI Topology” on page 583](#)
- [“Example NAT-mode ESI Topology” on page 588](#)
- [“Basic Regular Expression Syntax” on page 593](#)



Note: The ESI feature requires the Policy Enforcement Firewall Next Generation (PEFNG) license installed on the controller.

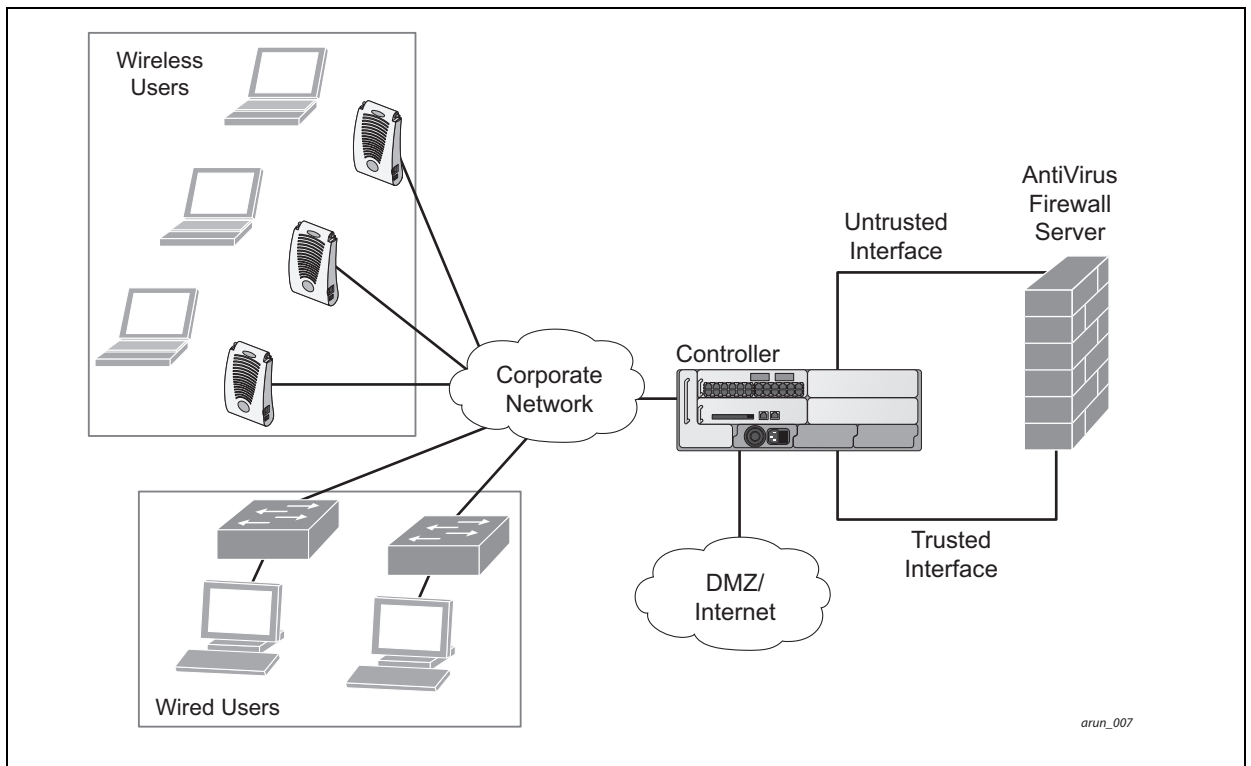
Understanding ESI

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.



Note: In ArubaOS 3.x, the only AVF server supported is Fortinet.

Figure 120 ESI-Fortinet Topology




In the topology in , 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 121](#) for an example).

Figure 121 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 122](#)) 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 573).

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 122 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 573). 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 (see “Peer Controllers” on page 573).

Figure 123 Peer Controllers



In this scenario, several controllers (master and local) are defined in the same syslog parser domain and are also configured to act as *peers*. From the standpoint of the ESI servers—because there is no good way of determining from which controller a given user came—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 “Basic Regular Expression Syntax” on page 593.) 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,” which 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.

ESI Configuration Overview

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.



Note: 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 Dell Management System, which is a suite of applications for monitoring multiple master controllers and their related local controllers and APs. Each application provides a Web-based user interface. The Dell Management System is available as an integrated appliance and as a software application that runs on a dedicated system. See the *Mobility Manager User Guide* for more information.

For more information about using these interfaces, see [Chapter 26, “Management Access”](#) .



Note: 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.



Note: 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. To apply the configuration (changes), click **Apply**. (The configuration will not take effect until you 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 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. To apply the configuration (changes), click **Apply**. (The configuration will not take effect until you click **Apply**.)

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. To apply the configuration (changes), click **Apply**. (The configuration will not take effect until you click **Apply**.)

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. To apply the configuration (changes), click **Apply**. (The configuration will not take effect until you click **Apply**.)
5. Refer to [Chapter 10, “Roles and Policies” on page 279](#), 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 (see).

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 [“Peer Controllers” on page 573](#)).

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 [on page 578](#)).
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 (based on the example shown in [Figure 123 on page 573](#)):

```
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 domains 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



Note: 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.



Note: 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 496).

4. Click **Apply** to commit the change, or click **Cancel** to discard the changes you made and exit the rule editing process.

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.



Note: 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
```

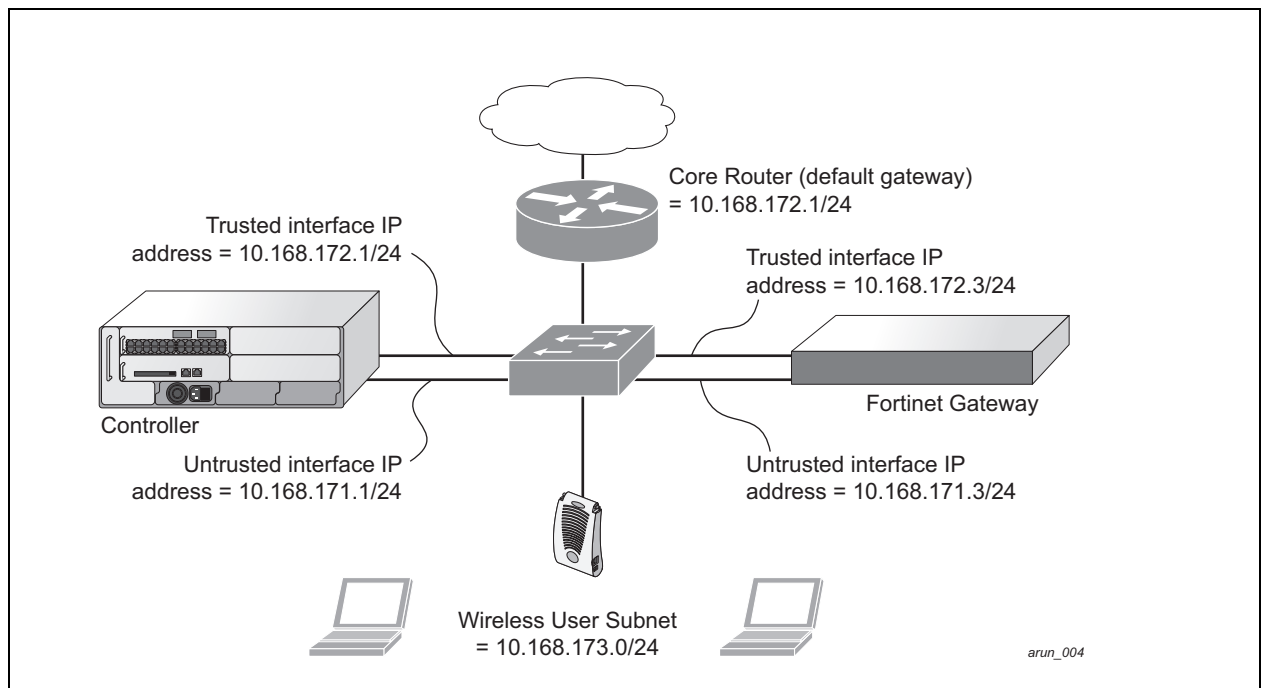
Example Route-mode ESI Topology

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. shows an example route-mode topology.



Note: ESI with Fortinet Anti-Virus gateways is supported only in route mode.

Figure 124 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.



Note: 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 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. To apply the configuration (changes), click **Apply**. (The configuration will not take effect until you 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. To apply the configuration (changes), click **Apply**. (The configuration will not take effect until you click **Apply**.)

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. To apply the configuration (changes), click **Apply**. (The configuration will not take effect until you click **Apply**.)

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. To apply the configuration (changes), click **Apply**. (The configuration will not take effect until you click **Apply**.)
4. Refer to [Chapter 10, “Roles and Policies” on page 279](#), 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.



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**.

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 124 on page 583](#).

```
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:

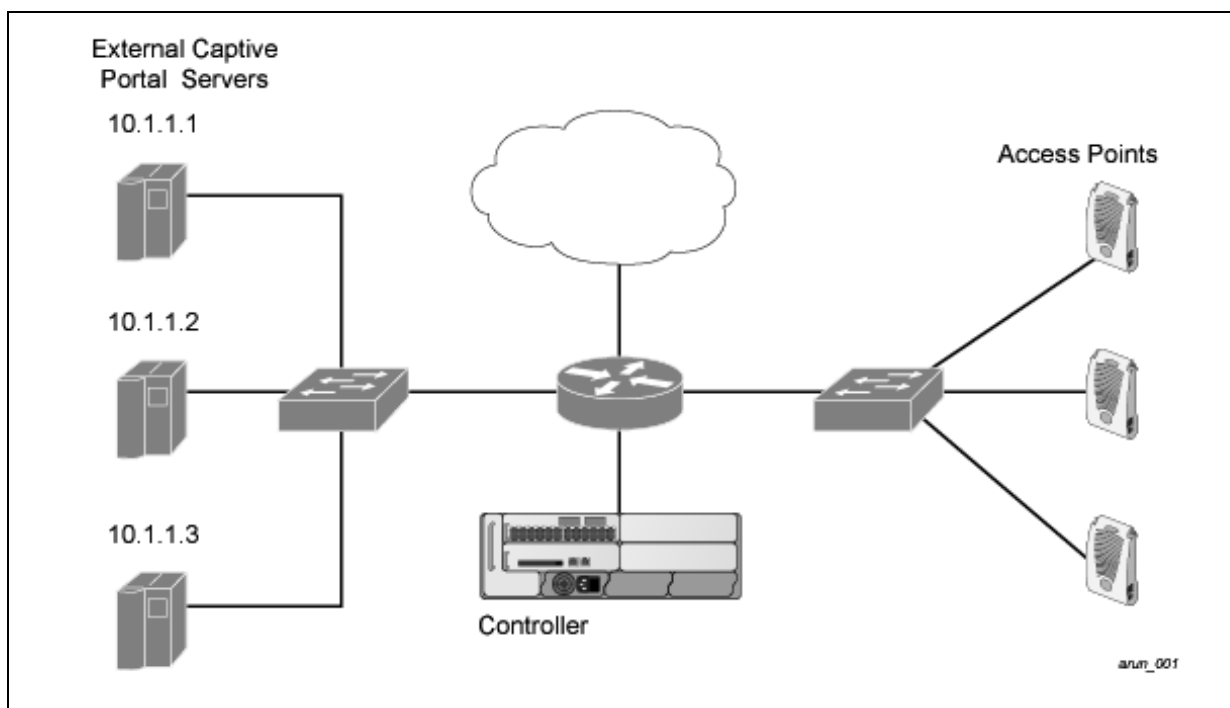
```
esi parser domain forti_domain
  server 10.168.172.3
```

```
esi parser rule forti_virus
  condition "log_id=[0-9]{10}[ ]"
  match ipaddr "src=(.*)" [ ]"
  set blacklist
  enable
```

Example 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 125 Example NAT-Mode Topology



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 [on page 575](#)).

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`.)



Note: 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.



Note: 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.



Note: 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. To apply the configuration (changes), click **Apply**. (The configuration will not take effect until you click **Apply**.)

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. on page 577](#)).

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. To apply the configuration (changes), click **Apply**. (The configuration will not take effect until you click **Apply**.)

Configuring the Example NAT-mode Topology in the CLI

The CLI configuration process consists of these general tasks:

- Configuring captive portal (see [Chapter 12, “Captive Portal”](#) 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

```

Basic Regular Expression 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 126](#)), repetition operators (described in [Table 127](#)), or expression anchors (described in [Table 128](#))—used to define 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 126 *Character-matching operators in regular expressions*

Operator	Description	Sample	Result
.	Match any one character.	grep .ord sample.txt	Matches <i>ford, lord, 2ord</i> , etc. in the file sample.txt.
[]	Match any one character listed between the brackets	grep [cng]ord sample.txt	Matches only <i>cord, nord</i> , and <i>gord</i>
[^]	Match any one character not listed between the brackets	grep [^cn]ord sample.txt	Matches <i>lord, 2ord</i> , etc., but not <i>cord</i> or <i>nord</i>
		grep [a-zA-Z]ord sample.txt	Matches <i>aord, bord, Aord, Bord</i> , etc.
		grep [^0-9]ord sample.txt	Matches <i>Aord, aord</i> , etc., but not <i>2ord</i> , 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 127](#) to search for multiple characters.

Table 127 *Regular expression repetition operators*

Operator	Description	Sample	Result
?	Match any character one time if it exists	egrep "?erd" sample.txt	Matches <i>berd, herd</i> , etc., <i>erd</i>
*	Match declared element multiple times if it exists	egrep "n.*rd" sample.txt	Matches <i>nerd, nrd, need</i> , etc.
+	Match declared element one or more times	egrep "[n]+erd" sample.txt	Matches <i>nerd, nnerd</i> , etc., but not <i>erd</i>
{n}	Match declared element exactly <i>n</i> times	egrep "[a-z]{2}erd" sample.txt	Matches <i>cherd, blerd</i> , etc., but not <i>nerd, erd, buzzerd</i> , etc.
{n,}	Match declared element at least <i>n</i> times	egrep ".{2,}erd" sample.txt	Matches <i>cherd</i> and <i>buzzerd</i> , but not <i>nerd</i>
{n,N}	Match declared element at least <i>n</i> times, but not more than <i>N</i> times	egrep "n[e]{1,2}rd" sample.txt	Matches <i>nerd</i> and <i>need</i>

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 128 Regular expression anchors

Operator	Description	Sample	Result
<code>^</code>	Match at the beginning of a line	<code>s/^/blah /</code>	Inserts "blah" at the beginning of the line
<code>\$</code>	Match at the end of a line	<code>s/\$/ blah/</code>	Inserts " blah" at the end of the line
<code>\<</code>	Match at the beginning of a word	<code>s/\</blah/</code>	Inserts "blah" at the beginning of the word
		<code>egrep "\<blah" sample.txt</code>	Matches <i>blahfield</i> , etc.
<code>\></code>	Match at the end of a word	<code>s/\>/blah/</code>	Inserts "blah" at the end of the word
		<code>egrep "\>blah" sample.txt</code>	Matches <i>souptlah</i> , etc.
<code>\b</code>	Match at the beginning or end of a word	<code>egrep "\bblah" sample.txt</code>	Matches <i>blahcake</i> and <i>countblah</i>
<code>\B</code>	Match in the middle of a word	<code>egrep "\Bblah" sample.txt</code>	Matches <i>sublahper</i> , 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: http://en.wikipedia.org/wiki/Regular_expression
- Regex syntax summary: <http://www.greenend.org.uk/rjk/2002/06/regexp.html>
- Basic regular expression (BRE) syntax: <http://builder.com.com/5100-6372-1050915.html>

A standards-compliant DHCP server can be configured to return the host Dell controller's IP address through the Vendor-Specific Option Code (option 43) in the DHCP reply. In the Dell user-centric network, this information can allow a Dell AP to automatically discover the IP address of a master controller for its configuration and management. This appendix describes how to configure vendor-specific option 43 on various DHCP servers.

This appendix contains the following topics:

- “Overview” on page 597
- “Windows-Based DHCP Server” on page 597
- “Linux DHCP Servers” on page 599

Overview

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.

Here is how option 43 works:

1. The DHCP client on a Dell AP adds an optional piece of information called the Vendor Class Identifier Code (option 60) to its DHCP request. The value of this code is **ArubaAP**.
2. The DHCP server sees the Vendor Class Identifier Code in the request and checks to see if it has option 43 configured. If it does, it sends the Vendor-Specific Option Code (option 43) to the client. The value of this option is the loopback address of the Dell master controller.
3. The AP receives a response from the DHCP server and checks if option 43 is returned. If it is, the AP contacts the master controller using the supplied IP address.

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

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 > Administrative Tools > 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 the **Add** button.
4. In the Option Type dialog box, enter the following information

Table 129 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 : ArubaAP
8. Click **OK** to save this information.
9. Under the server, select the scope you want to configure and expand it. Select **Scope Options** and expand it. 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 **ArubaAP** and click **OK**.
11. Confirm that the option **060 Dell Access Point** is listed in the right pane.

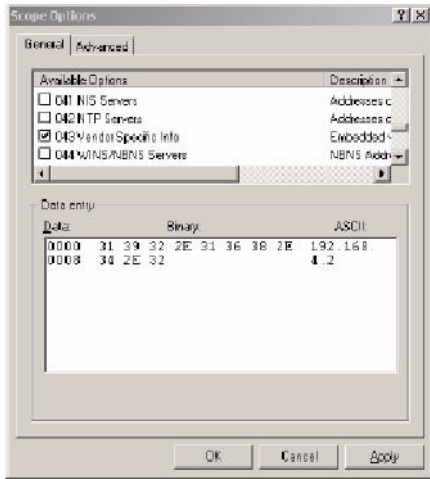
Configuring Option 43

Option 43 returns the IP address of the Dell master controller to an 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 126](#)), scroll down and select **043 Vendor Specific Info**

Figure 126 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 button 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 127 DHCP Scope Values



Linux DHCP Servers

The following is an example configuration for the Linux dhcpd.conf file.



Note: 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;
}
```

In many deployment scenarios, an external firewall is situated between Dell devices. This appendix 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. Note, however, that this appendix does not describe requirements for allowing specific types of user traffic on the network.



Note: 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.

This appendix includes the following topics:

- [“Communication Between Dell Devices” on page 601](#)
- [“Network Management Access” on page 602](#)
- [“Other Communications” on page 602](#)

Communication Between Dell Devices

This section describes the network ports that need to be configured on the firewall to allow proper operation of the Dell network.

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).

Between an AP and the master 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.

From an AP to the LMS controller:

- FTP (TCP port 21).
- TFTP (UDP port 69) for AP-52. For all other APs, if there is no local image on the AP (for example, a brand new AP) the AP will use TFTP to retrieve the initial image.
- NTP (UDP port 123).
- SYSLOG (UDP port 514).
- PAPI (UDP port 8211).
- GRE (protocol 47).

Between a Remote AP (IPsec) and a controller:

- NAT-T (UDP port 4500).
- TFTP (UDP port 69) .



Note: 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.

Network Management 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).

Other Communications

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. If the ArubaOS version is earlier than 2.5, allow SNMP traffic between the network management system and APs.
- 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.

This appendix contains the following topics:

- “Mode Support” on page 603
- “Basic System Defaults” on page 604
- “Default Management User Roles” on page 610
- “Default Open Ports” on page 613

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 130](#) describes the features that are not supported in each forwarding mode.

Table 130 *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 VLAN Pooling 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

Table 130 *Features not Supported in Each Forwarding Mode*

Forwarding Mode	Feature Not Supported
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 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

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 131](#) lists the predefined network services and their protocols and ports.

Table 131 *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

Table 131 *Predefined Network Services (Continued)*

Name	Protocol	Port(s)
svc-sip-tcp	tcp	5060
svc-kerberos	udp	88
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
svc-http-proxy2	tcp	8080
svc-http-proxy3	tcp	8888
svc-sips	tcp	5061
svc-v6-dhcp	udp	546 547

Table 131 *Predefined Network Services (Continued)*

Name	Protocol	Port(s)
svc-v6-icmp	icmp	0
any	any	0

Policies

The following are predefined policies.

Table 132 *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. <ol style="list-style-type: none"> Any HTTPS traffic destined for the controller will be NATed to port 8081, where the captive portal server will answer. All HTTP traffic to any destination will be NATed to the controller on port 8080, where an HTTP redirect will be issued. All HTTPS traffic to any destination will be NATed to the controller on port 8081, where an HTTP redirect will be issued. 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.

Table 132 *Predefined Policies (Continued)*

Predefined Policy	Description
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.
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.
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.

Table 132 *Predefined Policies (Continued)*

Predefined Policy	Description
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).
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.

Roles

The following are predefined roles.

Table 133 *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.

Table 133 *Predefined Roles (Continued)*

Predefined Role	Description
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 ipv6 session-acl v6-https-acl ipv6 session-acl v6-dhcp-acl ipv6 session-acl v6-icmp-acl ipv6 session-acl v6-dns-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.
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.

Table 133 *Predefined Roles (Continued)*

Predefined Role	Description
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 PEFNG 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.

Default Management User Roles

The ArubaOS software includes predefined management user roles.



Note: 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 134 *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: <code>http[s]://<ipaddress>[:port]/screens/wms/wms.login</code> . 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: <code>http[s]://<ipaddress>[:port]/screens/wms/wms.cgi?opcode=wlm-get-spot&campus-name=<campus id>&building-name<building id>&mac=<client1>,<client2>....</code>

Table 134 *Predefined Management Roles (Continued)*

Predefined Role	Permissions
network-operations	<p>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.</p> <p>This role permits the following WebUI pages and associated CLI commands:</p> <p>As a network-operations user, commands with an asterisk (*) are hidden in the CLI but are executed and visible from the WebUI.</p> <p>Plan Page</p> <ul style="list-style-type: none"> ● You can move APs on the floor plan and save their new location. ● You cannot change or modify the AP configuration. <p>Reports Page</p> <ul style="list-style-type: none"> ● You can view all of the available reports. <p>Events Page</p> <ul style="list-style-type: none"> ● You can view all of the available events. <p>Monitoring Page</p> <p>You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> ● show keys all ● show mobility-managers ● show roleinfo ● show mux config ● show mux stat ● show license* ● show ap essid ● DB:opcode=cr-load <p>Monitoring > Network > Network Summary</p> <p>You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> ● 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 mux state ● 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

Table 134 *Predefined Management Roles (Continued)*

Predefined Role	Permissions
<p>network-operations (continued)</p>	<p>Monitoring > Network > All Access Points Monitoring > Network > All Wired Access Points You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> ● 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 <p>Monitoring > Controller > Controller Summary You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> ● show switches ● show switches summary <p>Monitoring > Controller > Air Monitors You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> ● show wlan-ap start* <p>Monitoring > Controller > Clients You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> ● 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> <p>Monitoring > Controller > Clients > Packet Capture Monitoring > Controller > Clients > Locate Monitoring > Controller > Clients > Debug You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> ● aaa user debug mac <p>Monitoring > Controller > Clients > Disconnect You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> ● stm kick-off-sta <macaddr> ● aaa user logout <ipaddr>

Table 134 *Predefined Management Roles (Continued)*

Predefined Role	Permissions
network-operations (continued)	<p>Monitoring > Controller > Clients > Blacklist You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> • <code>stm add-blacklist-client <macaddr></code> • <code>aaa user delete {<ipaddr> all mac <macaddr> name <username> role <role>}</code> <p>Monitoring > Controller > Blacklist Clients You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> • <code>stm remove-blacklist-client <macaddr></code> <p>Monitoring > Controller > External Services Interface You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> • <code>show esi groups</code> • <code>show esi servers</code> • <code>show esi ping</code> • <code>show esi parser stats</code> <p>Monitoring > Controller > Ports You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> • <code>show model-switch-internal*</code> • <code>show slots</code> • <code>show private port status*</code> • <code>show vlan</code> <p>Monitoring > Controller > Inventory You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> • <code>show keys</code> <p>Monitoring > WLAN You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> • <code>DB:opcode=get-permissions</code> • <code>DB:opcode=cr-load</code> • <code>show switches</code> • <code>show switches summary</code> <p>Monitoring > Voice You can view the reports created by the following CLI commands:</p> <ul style="list-style-type: none"> • <code>show ap association voip-only</code> • <code>show ap active voip-only</code> • <code>show voice call-counters</code> • <code>show voice client status</code> • <code>show voice call-quality</code> • <code>show voice call-density</code> • <code>show voice call-cdrs</code> • <code>show voice call-perf</code>

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 135](#).

Table 135 *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.

Table 135 *Default (Trusted) Open Ports (Continued)*

Port Number	Protocol	Where Used	Description
21	TCP	controller	FTP server for AP6X software download.
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.
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.
500	UDP	controller	ISAKMP
514	UDP	controller	Syslog
1701	UDP	controller	L2TP
1723	TCP	controller	PPTP
2300	TCP	controller	Internal terminal server opened by <code>telnet soe</code> command.
3306	TCP	controller	Remote wired MAC lookup.
4343	TCP	controller	HTTPS. A different port is used from 443 in order to not conflict with captive portal. 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
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.

Table 135 *Default (Trusted) Open Ports (Continued)*

Port Number	Protocol	Where Used	Description
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	Dell Discovery Protocol (ADP).
8211	UDP	controller	For internal use.
8888	TCP	controller	Used for HTTP access.

This appendix 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 [Chapter 9, “802.1x Authentication”](#) 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 www.microsoft.com/downloads). Additional information on client configuration is available at <http://www.microsoft.com/technet/prodtechnol/winxppro/maintain/wificomp.msp#EQGAC>.

This chapter describes the following topics:

- “Configuring Microsoft IAS” on page 617
- “Configure Management Authentication using IAS” on page 621
- “Window XP Wireless Client Example Configuration” on page 624

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.




Note: The steps to perform this task may vary 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 www.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**.

Figure 128 IAS 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 [“Authentication with an 802.1x RADIUS Server” on page 262](#):

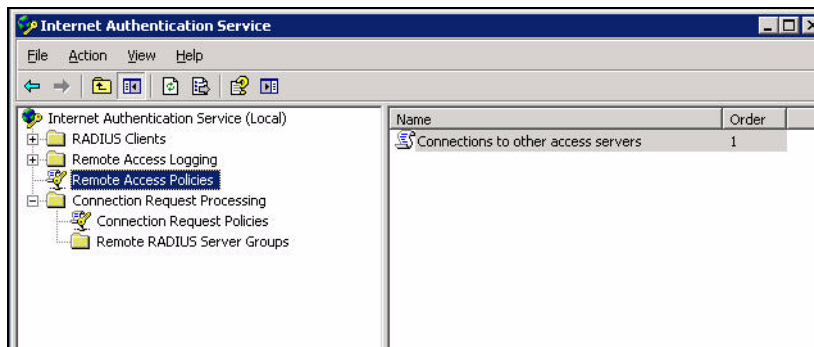
- The **Wireless-Computers** policy matches the “Domain Computers” group. This group contains the list of all computers that are members of the domain. This group is used for all computers to authenticate to the network.
- The **Wireless-Student** policy matches the “Student” group. This group is used 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:

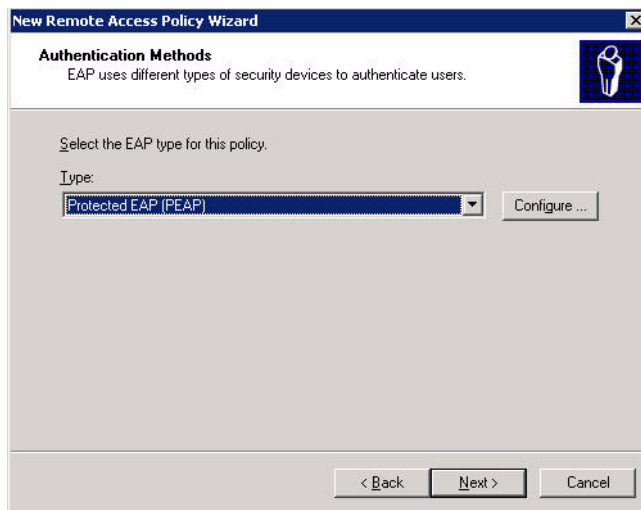
1. In the **Internet Authentication Service** window, select **Remote Access Policies**.

Figure 129 IAS 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.

Figure 130 Policy Configuration Wizard—Authentication Methods



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.

Figure 131 Policy Configuration Wizard—PEAP Properties



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”.

Figure 132 RADIUS class Attribute Configuration

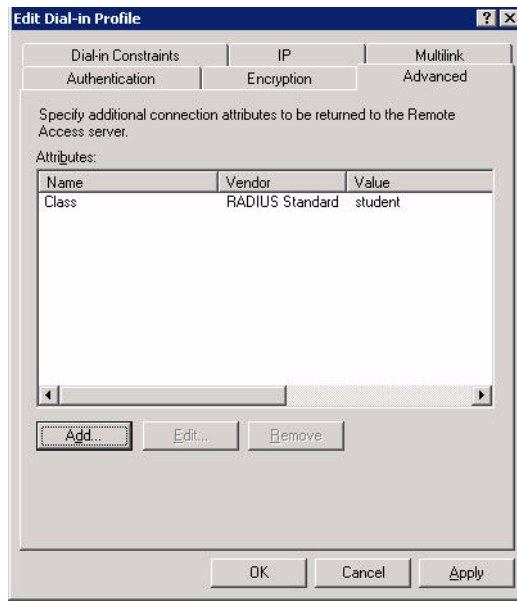


5. Click **OK**.

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.

Figure 133 Example RADIUS Class Attribute for “student”



Configure 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.



Note: The steps to perform this task may vary 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 www.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.

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.

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.

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.

Figure 134 Configuring a RADIUS Server for IAS Management Authentication

Security > Authentication > Servers

Servers | AAA Profiles | L2 Authentication | L3 Authentication | User Rules | Advanced

Server Group

- RADIUS Server
 - Management_RADIUS
 - RADIUS1
- LDAP Server
- Internal DB
- Tacacs Accounting Server
- TACACS Server
- XML API Server
- RFC 3576 Server
- Windows Server

RADIUS Server > Management_RADIUS

Show Reference | Save As | Reset

Host	127.0.0.1	Key Retype:
Auth Port	1812	Acct Port	1813
Retransmits	3	Timeout	5 sec
NAS ID		NAS IP	
Use MD5	<input type="checkbox"/>	Mode	<input checked="" type="checkbox"/>

- In the **Host** field, enter the IP address of the RADIUS server you want to use for Management Authentication.
- Enter and then retype the shared key for the server.
- Click **Apply**
- Select **Server Group** from the server list on the left window pane.
- In the entry blank on the right window pane, enter the name of a new server group (for example, “Management_group”), then click **Add**.
- Click **Apply**.
- Select the server group you just created from the list of server groups in the left window pane.
- In the Servers section, click **New**.
- Click the **Server Name** drop-down list and select your RADIUS server.

Figure 135 Configuring a Server Group for IAS Management Authentication

Security > Authentication > Servers

Servers | AAA Profiles | L2 Authentication | L3 Authentication | User Rules | Advanced

Server Group

- default
- group
- internal
- Management_group
- RADIUS Server
 - Management_RADIUS
 - RADIUS1
- LDAP Server
 - default
 - ldap
- Internal DB
- Tacacs Accounting Server
- TACACS Server

Server Group > Management_group

Show Reference | Save As | Reset

Fail Through

Servers

Name	Server Name	Server-Type	trim-FQDN	Match-Type	Operator	Match String	Actions
Internal (Local)	Internal (Local)		<input type="checkbox"/>	Authstring	contains		Add Rule Delete Rule
Management_RADIUS (Radius)							
RADIUS1 (Radius)							
default (LDAP)							
ldap (LDAP)							
tacacs (TACACS)							

Add Server | Cancel

Server Rules

Priority	Attribute	Operation	Operand	Type	Action	Value	Validated	Actions
New								

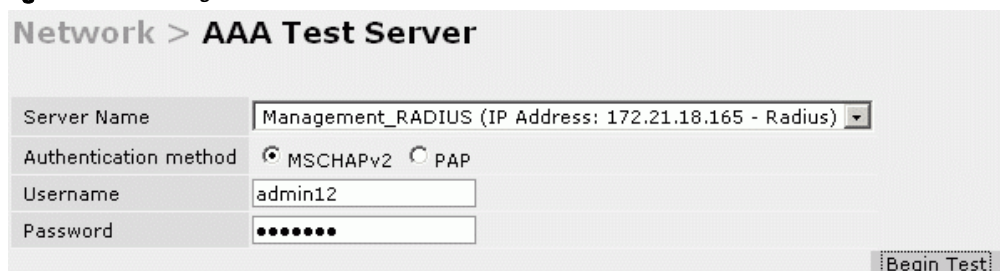
- Click **Apply**.

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.

Figure 136 Testing a RADIUS Server



Network > AAA Test Server

Server Name	Management_RADIUS (IP Address: 172.21.18.165 - Radius)
Authentication method	<input checked="" type="radio"/> MSCHAPv2 <input type="radio"/> PAP
Username	admin12
Password	••••••

Begin Test

Window XP Wireless Client Example Configuration


This section shows an example of how to configure a Windows XP wireless client using Windows XP's Wireless Zero Configuration service.



Note: 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.

Figure 137 Wireless Networks



4. Click the **Advanced** button to display the Networks to access window.

Figure 138 *Networks to Access*



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.

Note: 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 preshared 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 preshared key



Note: Do *not* select the option “This is a computer-to-computer (ad hoc) network; wireless access points are not used”.

Figure 139 shows the configuration for the SSID WLAN-01 which uses WPA network authentication with TKIP data encryption.

Figure 139 *Wireless Network Association*

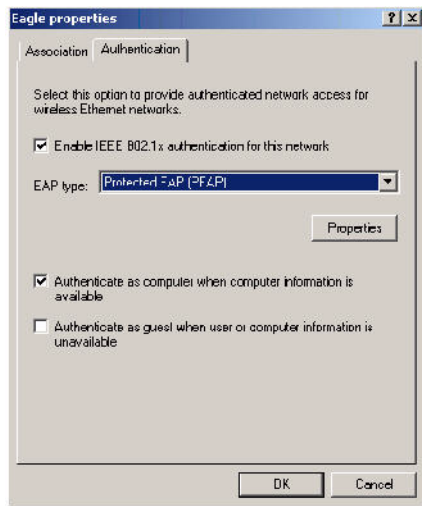


7. Click the **Authentication** tab to enter the 802.1x authentication parameters for the SSID. This tab configures the EAP type used between the wireless client and the authentication server.

Configure the following, as shown in Figure 140:

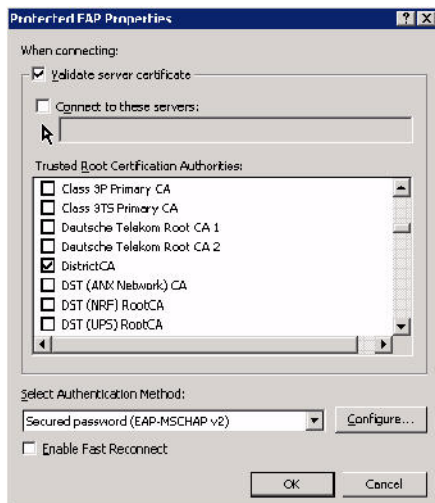
- Select Enable IEEE 802.1x authentication for this network.
- Select Protected EAP (PEAP) for the EAP type.
- Select Authenticate as computer when computer information is available. The client will perform computer authentication when a user is not logged in.
- Do not select Authenticate as guest when user or computer information is unavailable. The client will not attempt to authenticate as a guest.

Figure 140 *Wireless Network Authentication*




8. Under EAP type, select **Properties** to display the Protected EAP Properties window. Configure the client PEAP properties, as shown in [Figure 141](#):
 - 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.

Figure 141 *Protected EAP Properties*



9. 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.

Figure 142 *EAP MSCHAPv2 Properties*



You can customize the default captive portal page through the WebUI, as detailed in [Chapter 12, “Captive Portal”](#). This appendix discusses creating and installing a new internal captive portal page and other customization.

- “Creating a New Internal Web Page” on page 629
- “Installing a New Captive Portal Page” on page 631
- “Displaying Authentication Error Message” on page 631
- “Reverting to the Default Captive Portal” on page 632
- “Language Customization” on page 632
- “Customizing the Welcome Page” on page 635
- “Customizing the Pop-Up box” on page 636
- “Customizing the Logged Out Box” on page 637

Creating a New Internal Web 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 136](#):

Table 136 *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:

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:

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:

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, CSS files, 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.
- **Sygate Remediation Failure:** This is available as part of the External Services Interface feature and is outside the scope of this appendix.

Uploaded files can be referenced using:

```
https://<controller_IP>/upload/custom/<CP-Profile-Name>/<file>
```

Displaying Authentication Error Message

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. 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.

Language Customization

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](#), 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:

```

```

This should be replaced with a link like this:

```

```

- 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":
            localised_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.

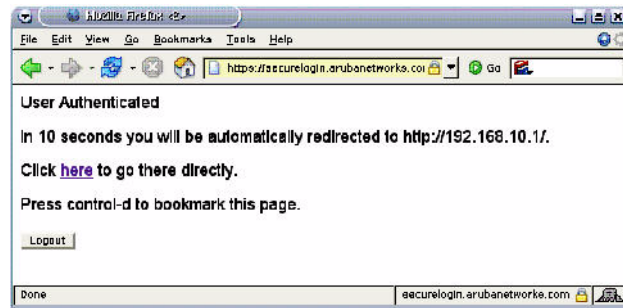


Note: The “Registered User” and “Guest User” sections of the login page are implemented as graphics files, referenced by the default CSS styles. In order to change these, you will need to create new graphic files, download the CSS file, edit the reference to the graphics files, change the style reference in your index file and then upload all files as “content” to the controller.

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 143](#):

Figure 143 *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 Pages** and 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 143](#) 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>
```

This section of the document provides instructions and information on using VIA.

Pre-requisites

Ensure that the end-user system meets the following pre-requisites:

- VIA can be installed only on systems running:
 - Microsoft Windows XP with SP2
 - Microsoft Windows Vista
 - Microsoft Windows 7
- Requires the following Microsoft KB on the end-user systems:
 - On Microsoft Windows XP SP2—KB918997 (<http://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 (<http://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 [Figure 58 on page 338](#).

In this example, they can download VIA set up files from <https://115.52.100.10/via> after entering their corporate credentials.

Figure 144 Login to Download VIA

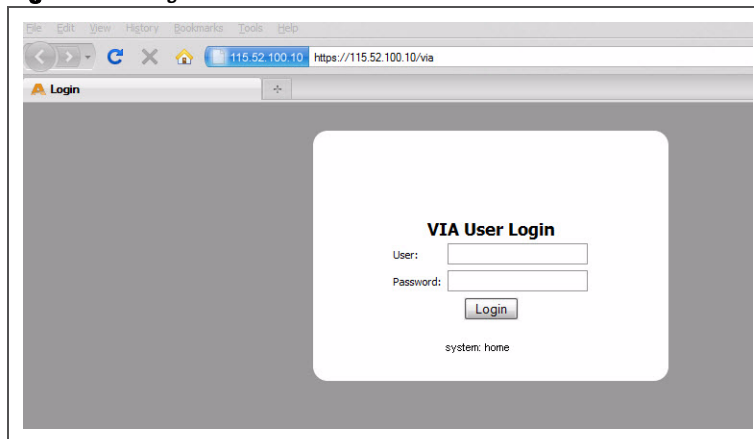



Figure 145 Downloading VIA set up file after authentication



Installing VIA

Double click the downloaded set up file (*ansetup.msi*) to start the installation process. Ensure that you have met the pre-requisites before proceeding with the installation.

Using VIA

The VIA 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 this 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.

Diagnostics Tools

- **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.
- **Aruba VIA info**—Information about 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.

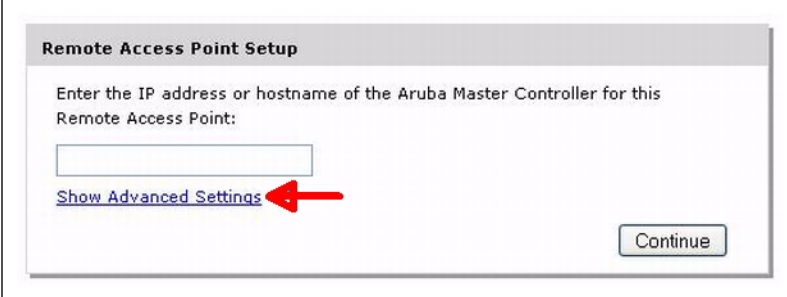
- **Log Settings**—Allows you to set VIA log levels. By default, the log level is set to *Trace*. This setting captures extensive activity information about VIA.
- **Connection Profile**—Allows you to select and connect to a different connection profile. This is usually useful if you are in remote location 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.

This document provides information on provisioning your remote AP (RAP) at home using a static IP address, PPPoE connection, or 3G/EVDO USB modem.

You provision the RAP using provisioning wizard:

1. Navigate to the RAP configuration URL — .
2. Enter the IP address or hostname of the controller.
3. Click the **Show Advanced Settings** link, shown in [Figure 146 on page 643](#).

Figure 146 *Show Advanced Settings*



The screenshot shows a web form titled "Remote Access Point Setup". The form contains the following text: "Enter the IP address or hostname of the Aruba Master Controller for this Remote Access Point:". Below this text is an empty text input field. Underneath the input field is a blue hyperlink labeled "Show Advanced Settings" with a red arrow pointing to it from the right. In the bottom right corner of the form is a "Continue" button.

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.

Provision the RAP using a Static IP Address

Select the **Static IP** tab and enter the required details. See [Table 137](#) for information on parameters.

Figure 147 Provision RAP using Static IP

Static IP | PPPoE | USB

IP Address

Netmask

Gateway

Primary DNS

Domain

Save Clear

Continue

Table 137 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.
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 the Save button 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 138](#) for information on parameters

Figure 148 Provision RAP on a PPPoE Connection

The screenshot shows a web-based configuration interface for provisioning a RAP. At the top, there are three tabs: 'Static IP', 'PPPoE', and 'USB'. The 'PPPoE' tab is currently selected. Below the tabs, there are three text input fields labeled 'Service name', 'Username', and 'Password'. Below these fields are two buttons: 'Save' and 'Clear'. At the bottom right of the main form area, there is a 'Continue' button.

Table 138 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 the Save button after you have entered all the details.

Using 3G/EVDO USB Modem

The following procedure illustrates provisioning your RAP using a 3G/EVDO USB modem.

Select **USB** tab and select your modem from the drop down list. For some common modems, the details are automatically filled.

Figure 149 Provision using a pre-configured USB Modem

The screenshot shows a configuration window with three tabs: Static IP, PPPoE, and USB. The USB tab is active. The 'Device' dropdown is set to 'Other (Any)'. A dropdown menu is open for 'Device Type', listing several modem models: 'Other (Any)', 'USBConnect 881 (ATT)', 'USB 598 / U597 / Compass 597 (Sprint/Verizon)', 'Ovation U727 / U720 / U300 (Sprint/Verizon)', 'UM175 / UM150 (Verizon)', 'Mercury Sierra Compass 885 (ATT)', and 'Quicksilver Globetrotter ICON 322 (ATT)'. Below the dropdowns are input fields for 'Initialization String', 'PPP Username', 'PPP Password', 'TTY Device Path', 'Device Identifier', and 'Dial String'. At the bottom are 'Link Priority Cellular' and 'Link Priority Ethernet' fields, both set to '0'. There are 'Save' and 'Clear' buttons at the bottom of the main form, and a 'Save' button at the bottom right of the window.

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 150 Provision using a USB Modem with Custom Settings

The screenshot shows the same configuration window as Figure 149, but with 'Device Type' set to 'any'. The 'Initialization String', 'PPP Username', 'PPP Password', 'TTY Device Path', 'Device Identifier', and 'Dial String' fields are empty. The 'Link Priority Cellular' and 'Link Priority Ethernet' fields are set to '0'. There are 'Save' and 'Clear' buttons at the bottom of the main form, and a 'Continue' button at the bottom right of the window.

- 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.

Click the **Save** button after you have entered all the details and click the **Continue** button to complete provisioning of your RAP.

Numerics

- 20 MHz channel assignment 133
- 40 MHz channel assignment 133
- 802.11n zone 86
- 802.1x authentication
 - configuring 251

A

- AC
 - mappings 562
 - types 561
- access category. *See* AC
- access control lists 280
- Access Points
 - deploying 49
 - high-latency link deployments 135
 - IP addresses 50
 - low-speed deployments 135
- accounting
 - configuring 247
- ACL white list 282
- ACLs and remote APs 180
- air monitoring and mesh 224
- AP
 - configuring 101
 - status
 - down 83
 - up, live 83
- AP failback 136
- AP groups 103
- AP installation modes 132
- AP maintenance mode 137
- area
 - 802.11n zone 86
 - don't care 86
 - don't deploy 86
- ARM 139
 - ARM metrics 150
 - band steering 147
 - spectrum load balancing 149
 - traffic shaping 148
 - troubleshooting 151
- authentication 516
- authentication methods
 - smart card 516
 - static 516
 - username and password 516

- authentication server
 - configuring timers 249
 - trim domain information 243
- authentication server group
 - configuring 231
 - configuring rules 244
 - fail-through 240
 - FQDN server selection 241
 - order of servers 240
 - server selection 241

B

- backhaul, wireless 195
- backup configuration, remote APs 169
- basic regular expression syntax 593
- blacklisting clients 484

C

- campus AP whitelist 376
- captive portal 520
 - changing to HTTP protocol 317
 - configuring 299
 - default page customization 320
 - different VLAN clients 319
 - per-SSID configuration 316
 - proxy Web server configuration 318
- captive portal page
 - customizing 320
- care-of address 399
- certificates 503
 - AAA FastConnect 260
 - importing 505
 - obtaining server certificate 504
 - SSH access 494
 - WebUI management 493
- channel assignment, 20 MHz 133
- channel assignment, 40 MHz 133
- Channel Reuse 149
- channel switch announcement 133
- client blacklisting 484
- cluster profile, mesh
 - overview 193
- Configuring WISPr authentication 296
- connecting to network 49
- control plane security 373

D

- dead peer detection
 - configuring 366
- deployment considerations, mesh 198
- DHCP client 61
- DHCP with option 43 597
- dialer
 - configuring 366
- don't care 86
- don't deploy 86
- double encryption 166
- duplicate AP names 103
- duplicate names 103

E

- enable mode password reset 499
- example configuration
 - 802.1x 262
 - captive portal 305
- example configurations
 - WLANs 115
- External Services Interface
 - configuring 569
 - syslog parser 571

F

- failback, remote APs 178
- file transfer 521
- firewall parameters 289
- flash backup and restore 522
- floor
 - 802.11n zone 86
 - don't care 86
 - don't deploy 86
- foreign agent 399
- foreign network 399
- Fortinet topology 570

G

- GRE tunnel
 - configuring 67
- guest access pass 517
- guest accounts 518
- guest provisioning 510
 - guest accounts 518
 - print guest account information 519

H

- high-throughput, virtual AP profile 125
- home agent 399
- home agent table 400
- home network 399

I

- IDS
 - configuring 463
- image file transfer 522
- indoor AP 132
- initial setup 44
- Internal AP 428
- internal database
 - configuring 236
- IP mobility 399
- IPv6 535

L

- L2TP
 - configuring 349
- LACP
 - Best Practices 491
 - configuring 489
 - configuring with WebUI 491
 - data units (DUs) 489
 - sample configuration 492
 - Tx/Rx 489
 - with the CLI 489

LAG

- group 489
- member ports 489
- LDAP server
 - configuring 234
- Link Aggregation Control Protocol
 - see LACP 489
- Link Aggregation Group
 - see LAG 489
- log files, copying 523
- logging
 - configuring 509
- loopback address
 - configuring 67
- loopback IP address 66

M

- MAC-based authentication
 - configuring 369
- maintenance mode, AP 137
- management authentication
 - configuring 246
- mesh
 - bridging 219
 - deployment considerations 198
 - secure jack 220
 - statistics 225
 - troubleshooting 221
 - tunneling 220
 - wired AP profile 219

- mesh cluster 190
- mesh link
 - creating 191
 - overview 191
- mesh nodes, provisioning 190, 222
- mesh path 190
- mesh point
 - behavior 190
 - boot sequence 224
 - overview 190
- mesh portal
 - behavior 190
 - boot sequence 224
 - overview 190
- mesh service set identifier. *See* *MSSID*
- migration 421
- mobile client 399
- mobility domain 399
 - configuring 400
 - example configuration 402
- Mounting Devices 439
- MP. *See* *mesh point*
- MPP. *See* *mesh portal*
- MSSID 190
- Multi-function Media Eject Button 440

N

- Network-attached storage 437
- Network-Attached Storage (NAS) 437
- NTP
 - configuring 524

O

- option 43 on DHCP server 597
- outdoor AP 132

P

- password recovery 499
- policies 279
 - configuring 280
- port
 - configuring 57
- PPPoE client 61
- PPTP
 - configuring 362
- presared key 393
- print guest account information 519
- Print Server 442
- profiles
 - configuring 105
- profiles, mesh
 - cluster 193
 - recovery 194

- provisioning
 - mesh caveats 223
 - mesh nodes 190, 222
 - outdoor APs 190, 222
 - remote APs 161

PSK 393

Q

- QoS for voice
 - configuring 545

R

- radio profile, mesh
 - configuring 199
 - parameters 128, 129, 199, 202, 550, 566
- RADIUS server
 - configuring 232
- RADIUS Server Authentication Codes 233
- RAP Local Network Access 179
- RAP Static Inner IP Address 237
- recovering password 499
- recovery profile, mesh 194
- remote AP
 - ACLs 180
 - backup configuration 169
 - configuring 153
 - DNS setting 177
 - failback 178
 - provisioning 161
 - split tunneling 180
 - WMM 186
- removing duplicate AP names
 - AP names 103
- restrict to one guest 520
- RF Plan 49, 71
 - add background image, name first floor 98
 - add background image, name second floor 98
 - add/edit floors 97
 - coverage maps, heat maps 83
 - create a building 96
 - create area
 - don't care 98
 - don't deploy 99
 - down AP icon 83
 - exporting 90
 - image guidelines 84
 - importing 90
 - model access points 97
 - model air monitors 97
 - run RF Plan 99
 - run the AM plan 100
 - up AP icon 83

- role
 - assigning 286
 - configuring 283
- route-mode topology 583

S

- secure jack and mesh 220
- server derivation rules
 - configuring 289
- server group
 - assigning 246
 - configuring 231, 239
- server rules
 - configuring 244
- server-derived role 286
- site-to-site VPN
 - configuring 363
- smart card authentication 516
- SNMP
 - configuring 507
- solutions, mesh
 - overview 195
 - point-to-multi-point 196
 - point-to-point 196
 - wireless backhaul 195
 - with thin APs 195
- source NAT 63
- source NAT and dynamic VLAN 63
- split tunneling, remote APs 180
- Stateful authentication 293
- static authentication method 516
- static route
 - configuring 65
- static routes 65
- syslog parser 571

T

- TACACS+ server
 - configuring 235
- timers
 - authentication 249
- tunnel, GRE 67

U

- Uplink Manager 429
- USB Cellular Modem 428
- USB Modem 433
 - configuring 433
- user derivation rules
 - configuring 287
- user role
 - assigning 286
 - configuring 283
- user-derived role 286
- username and password authentication 516

V

- virtual AP profile, high-throughput 125
- virtual APs 105
- VLAN
 - assignment 60
 - configuring 55
 - disabling VLAN routing 65
 - dynamic address 60
 - inter-VLAN routing 64
 - static address 60
- Voice Services Module
 - features 553
- VoIP
 - configuring for 545
- VPN AAA deployments 348
- VRRP
 - configuring 413
- VSA-derived role 289

W

- WebUI 37, 41, 45
- white list 282
- whitelist synchronization 381
- whitelisting ACLs 282
- Wi-Fi Multimedia. *See* WMM
- Windows authentication server 235
- wireless backhaul 195
- WISPr authentication 293
- wizard
 - AP 37, 41
 - license 37, 41, 531, 532
 - setup 44, 49
 - WLAN 37, 41
 - 102
- wizard, wlan 300
- WLAN policy configuration 475
- WMM
 - AC mapping 562
 - enabling 561
 - remote AP support 186

X

- xSec
 - configuring 323
 - configuring for wired clients 325
 - configuring for wireless clients 324
 - configuring wireless clients 327